arXiv:cs/0602040v1 [cs.LO] 10 Feb 2006

PLTL Partitioned Model Checking for Reactive Systems under Fairness Assumptions

S. Chouali, J. Julliand, P.-A. Masson and F. Bellegarde
Laboratoire d'Informatique de l'Université de Franche-Comté,
FRANCE - LIFC FRE CNRS 2661

ACM Transactions on Embedded Computing Systems (TECS), 4(2):267–301, May 2005

Abstract

We are interested in verifying dynamic properties of finite state reactive systems under fairness assumptions by model checking. The systems we want to verify are specified through a top-down refinement process.

In order to deal with the state explosion problem, we have proposed in previous works to partition the reachability graph, and to perform the verification on each part separately. Moreover, we have defined a class, called $\mathcal{B}_{mod}$, of dynamic properties that are *verifiable by parts*, whatever the partition. We decide if a property P belongs to $\mathcal{B}_{mod}$ by looking at the form of the Büchi automaton that accepts $\neg P$. However, when a property P belongs to $\mathcal{B}_{mod}$, the property $f \Rightarrow P$, where f is a fairness assumption, does not necessarily belong to $\mathcal{B}_{mod}$.

In this paper, we propose to use the refinement process in order to build the parts on which the verification has to be performed. We then show that with such a partition, if a property P is verifiable by parts and if f is the expression of the fairness assumptions on a system, then the property $f \Rightarrow P$ is still verifiable by parts.

This approach is illustrated by its application to the chip card protocol T=1 using the B engineering design language.

keywords. Refinement design, PLTL model checking, fairness assumptions, out-of-core model checking.

1 Motivations

This paper is about the verification of dynamic properties of finite state systems. In our approach, reactive systems are modeled by transition systems expressed as event systems, for example in B [Abr96], and are specified through a top-down refinement process. We propose to express the dynamic properties (safety, liveness) as formulae of the Propositional Linear Temporal Logic (PLTL) [MP92, AM98], and to verify them by model checking [QS82, CES86, CGP99].

It is well known that the main drawback of the PLTL model checking [LP85, VW86] is that it cannot handle very large finite systems. To avoid the model checking explosion, many solutions have been proposed, such as partial order [KP88, WG93], abstraction [CC77, CGL94, DF95] and symbolic representation by BDD [BCMD90, McM93]. For a class of PLTL properties, we propose another solution which is compatible with the previous ones.

1.1 Our Propositions

In order to deal with the problem of the exponential blow up of the PLTL model checking, we have proposed in [JMM01, MMJ00] a method that relies on a partitioning of the transition system into several parts. The

properties are verified on each part separately by an out-of-core [Tol99] model checking technique. As every part is verified separately from the others, the other parts can be stored on disks while the part of interest is in the main memory. We call *verifiable by parts* the properties that are such that if they hold on every part of a transition system (whatever the partitioning), then they hold on the whole transition system. A sufficient condition C on the Büchi automaton which accepts the ω-language of the negation of a property $\neg P$ allows us to decide if P is verifiable by parts. C is a syntactic condition on the Büchi automata. Safety and liveness properties such as $\Box(p \Rightarrow \bigcirc q)$, $\Box(p \Rightarrow \Diamond q)$ and $\Box(p \Rightarrow r\mathcal{U}q)$ are verifiable by parts.

The fact that a property P is verifiable by parts does not depend on the way the parts are constructed. But, when verifying P by parts, the fact that P holds on *every* part does depend on it. We have proposed a partitioning based on the refinement process.

Notice that choosing a partitioning method allows more PLTL properties to be verifiable by parts. Some PLTL properties that are not verifiable by parts for all possible partitionings might become verifiable by parts under the hypothesis of this particular partitioning.

In this paper, we extend our verification method to liveness properties on transition systems provided with fairness assumptions. To verify a property P under the fairness assumptions f, it is necessary to verify that $f \Rightarrow P$ holds on the transition system. The problem for verifying $f \Rightarrow P$ by parts is that even if P is verifiable by parts for all partitionings, the Büchi automata of $\neg(f \Rightarrow P)$ does not in general satisfy the condition C. However, with the fair refinement based partitioning proposed in the paper, it is enough that a property P is verifiable by part for all partitionings, i.e. satisfy C, for having $f \Rightarrow P$ verifiable by this particular partitioning.

1.2 Related Works

The component verification of a multiprocess system is achieved by verifying the properties separately on each component. Then, the compositionality with respect to parallel composition ensures that the properties hold on the whole system (see [CLM89, KL93, KV98, AdAHM99, KKPR99]). Generally, this method allows verifying safety properties because a component extracted from its environment is an abstraction of the parallel system. For example, [KL93] proves that $C_1 \| C_2$ satisfies an invariant $I_1 \wedge I_2$ by proving that $I_2 \Rightarrow I_1$ holds on C_1, and that $I_1 \Rightarrow I_2$ holds on C_2. These methods are called *assume guarantee paradigm*. Some methods, like in [CGK97], use also a component verification of liveness properties. In our approach, a part is not a component of a parallel composition. Our partitioned method is *additive* whereas the compositionality is *multiplicative*.

In this paper we present a fair refinement that can be compared with the works on fair simulation presented in [HKR97, GL94, KV96, DHWT91]. The fair refinement differs from the fair simulation in the following points:

- The fair refinement concerns fair transition systems which models action systems, whereas the fair simulation is concerned by the "state systems" modeled as Kripke structures. In the fair refinement, fairness assumptions are expressed on transitions. In Kripke structure, fairness assumptions are expressed on states.

- The fair refinement is a fair τ-simulation because we add new actions in the refined system which are not observable (τ-actions) in the abstract one.

- The fair refinement is a state simulation, as is the fair simulation, but it is also an action simulation, i.e. the sequence of abstract actions is a τ-simulation of the sequence of refined actions. This means that the refined sequence of actions is identical to the sequence of abstract actions, in which finite sub-sequence of τ-actions are interleaved.

- The verification of the fair refinement is linear in the size of the refined system, but the verification of the fair simulation is polynomial [HKR97].

Our method can be combined with the parallel composition of components. We have proved in [BJK02] that the parallel composition is compatible with the refinement, i.e. if C_1 is refined by C_3 then $C_1 \parallel C_2$ is refined by $C_3 \parallel C_2$.

Paper Organization. The paper is organized as follows. The preliminaries (Section 2) presents the notation and the concepts of fair transition systems, PLTL and Büchi automata. After a presentation of our refinement relation in Section 3, we explain the partitioned model checking technique in Section 4. Section 5 studies the extension of the partitioned verification principles to fair transition systems. Section 6 gives an example of an application to this technique, and some experimental results are discussed in Section 7. Finally, we situate our concerns and give some ideas about future works in Section 8.

2 Preliminaries

2.1 Transition systems

We introduce a finite set of *variables* $x \in V$ with their respective finite domains D_x. Let $AP_V \stackrel{\text{def}}{=} \{x = v \mid x \in V,\ v \in D_x\}$ be a set of atomic propositions over V.

Definition 2.1 (Transition System) *Let Act be a nonempty alphabet of labels (names of actions). A transition system $TS \stackrel{def}{=} \langle S_0, S, Act, \rightarrow, L\rangle$ interpreted over V has a set of initial states S_0 included in a finite set of states S, a labelled transition relation $\rightarrow \subseteq S \times Act \times S$ that must be total, and a state labelling function $L : S \rightarrow 2^{AP_V}$.*

This is a labelled and interpreted transition system. A labelled transition relation $\rightarrow$ is a set of triples (s, a, s') (written as "$s \stackrel{a}{\rightarrow} s'$"). It is an interpreted transition system because each state is decorated with a set of atomic propositions. Notice that the set of atomic propositions which is associated to a state s must be consistent, i.e. if $x = v \in L(s)$ and $x = v' \in L(s)$ then $v = v'$.

Remark 2.2 *As the transition relation is total, there can be no deadlock in a transition system. If a state s has no successor, a transition $s \stackrel{\texttt{Skip}}{\rightarrow} s$ (where `Skip` does not belong to Act) is added to obtain a transition system.*

Definition 2.3 (Execution) *An execution of a transition system $\langle S_0, S, Act, \rightarrow, L\rangle$ is an infinite alternating sequence $\sigma \stackrel{def}{=} s_0 \stackrel{a_0}{\rightarrow} s_1 \stackrel{a_1}{\rightarrow} s_2 \cdots s_i \stackrel{a_i}{\rightarrow} s_{i+1} \cdots$ of states and actions such that $s_0 \in S_0$ and for every $i \geq 0$, we have $s_i \stackrel{a_i}{\rightarrow} s_{i+1} \in \rightarrow$ [KLM+02].*

We denote by Σ_{TS} the set of all the executions of a transition system TS.

Definition 2.4 (Fragment of an execution) *We say that σ' is a fragment of an execution σ (written as $\sigma' \subset \sigma$) if the sequence of transitions (finite or infinite) executed in σ' is also executed in σ.*

We note $s \stackrel{*}{\rightarrow} s'$ (resp. $s \stackrel{+}{\rightarrow} s'$) the fragments of executions leading from s to s' by executing zero (resp. one) or many transitions.

Definition 2.5 (Cycle) *A cycle of a transition system TS is a finite fragment $c \stackrel{def}{=} s_0 \stackrel{a_0}{\rightarrow} s_1 \stackrel{a_1}{\rightarrow} \cdots \stackrel{a_{n-1}}{\rightarrow} s_n$ of an execution of TS, such that $s_n = s_0$ and for all $0 \leq i, j < n$, if $i \neq j$ then $s_j \neq s_i$.*

We extend a finite fragment $\sigma' \stackrel{\text{def}}{=} s_i \stackrel{a_i}{\rightarrow} s_{i+1} \cdots s_{k-1} \stackrel{a_{k-1}}{\rightarrow} s_k$ to an execution $\sigma \stackrel{\text{def}}{=} s_i \stackrel{a_i}{\rightarrow} s_{i+1} \cdots s_{k-1} \stackrel{a_{k-1}}{\rightarrow} s_k \stackrel{\texttt{Skip}}{\rightarrow} s_k \stackrel{\texttt{Skip}}{\rightarrow} s_k \cdots$ and we call it an *extension*.

Definition 2.6 (Trace of (a fragment of) an execution) *The trace of an execution or a fragment of an execution σ, written as $tr(\sigma)$, is the sequence of the labels of the transitions executed in σ.*

Let TS_2 be a transition system. The τ-transition system of TS_2 on Act_1, written as τ-TS_2, is a transition system identical to TS_2 where the names of the actions of Act_2 which do not belong to Act_1 are named τ.

Definition 2.7 (τ-transition system) *Let Act_1 bet a set of actions. Let $TS_2 = \langle S_{0_2}, S_2, Act_2, \rightarrow_2, L_2\rangle$ be a transition system, such that $Act_1 \subseteq Act_2$. We call τ-transition system of TS_2 on Act_1, the transition system τ-$TS_2 = \langle S_{02}, S_2, Act_{1\tau}, \rightarrow_{2\tau}, L_2\rangle$ such that $Act_{1\tau} = Act_1 \cup \{\tau\}$, and the relation $\rightarrow_{2\tau}$ is defined as:*

$$\rightarrow_{2\tau} \stackrel{def}{=} \{s_2 \xrightarrow{\tau} s_2' \mid s_2 \xrightarrow{a_2} s_2' \in \rightarrow_2 \ \wedge\ a_2 \in Act_2 \wedge a_2 \notin Act_1\} \cup \{s_2 \xrightarrow{a_1} s_2' \mid s_2 \xrightarrow{a_1} s_2' \in \rightarrow_2 \ \wedge\ a_1 \in Act_1\}.$$

Any cycle c such that $tr(c) \in \tau^*$ is called a τ-cycle.

2.2 Fair transition systems

We note $Inf_s(\sigma)$ the set of states occurring infinitely often in an execution σ. We note $Inf_t(\sigma)$ the set of transitions occurring infinitely often in σ. We note $In(T)$ the set of states which are source of transitions in a set of transitions T. That is,

- $Inf_s(\sigma) \stackrel{\text{def}}{=} \{s \mid \sigma = s_0 \xrightarrow{a_0} s_1 \cdots s_i \xrightarrow{a_i} s_{i+1} \cdots \wedge\ s = s_i \text{ for infinitely many } i\}$,
- $Inf_t(\sigma) \stackrel{\text{def}}{=} \{\ t \mid \sigma = s_0 \xrightarrow{a_0} s_1 \cdots s_i \xrightarrow{a_i} s_{i+1} \cdots \wedge\ t = s_i \xrightarrow{a_i} s_{i+1} \text{ for infinitely many } i\}$,
- $In(T) \stackrel{\text{def}}{=} \{s \mid s \xrightarrow{a} s' \in T\}$.

Notice that in [MP95], fairness requirements are expressed on transitions. Here, we express fairness requirements on actions. Fair transition systems model fair action systems. Therefore, the user expresses fairness constraints in the form of fair actions. In a fair transition system, a fair action is expressed by a set F_i of the transitions which have the name of this action as a label.

Definition 2.8 (Fair transition system) *Let $TS = \langle S_0, S, Act, \rightarrow, L\rangle$ be a transition system. Let $\mathcal{F}$ be a set of fairness constraints $\{F_1, F_2, \cdots, F_m\}$ where every $F_i \subseteq \rightarrow$ is a set of the transitions expressing one fairness constraint. The fair transition system FTS is the tuple $\langle S_0, S, Act, \rightarrow, L, \mathcal{F}\rangle$ (often written as $\langle TS, \mathcal{F}\rangle$) such that*

(a) $(s_1 \xrightarrow{a} s_1' \in F_i \ \wedge\ s_2 \xrightarrow{b} s_2' \in F_i) \Rightarrow (a = b)$,

(b) $(s \xrightarrow{+} s' \in \rightarrow^+ \ \wedge\ s \xrightarrow{a} s' \in F_i) \Rightarrow \exists (s_1, s_1') \cdot (s \xrightarrow{+} s' = s \xrightarrow{*} s_1 \xrightarrow{a} s_1' \xrightarrow{*} s' \wedge s_1 \xrightarrow{a} s_1' \in F_i)$,

(c) $(s \xrightarrow{a} s' \in F_i \ \wedge\ s \xrightarrow{a} s'' \in \rightarrow) \Rightarrow (s' = s'')$.

Let us comment on the three constraints (a), (b) and (c).

(a) Each element F_i of $\mathcal{F}$ expresses one fairness assumption involving one fair action. Therefore, all the transitions of F_i have the same label.

(b) For each transition $s \xrightarrow{a} s'$ of F_i, there does not exist a fragment of execution beginning in s and ending in s' that does not contain a fair transition of F_i. This constraint allows expressing fairness as a PLTL property (see Section 2.4)

(c) We require determinism for a fair action since it comes from the environment. So, for each transition $s \xrightarrow{a} s'$ of F_i, there does not exist another transition $s \xrightarrow{a} s''$, such that $s' \neq s''$. This constraint allows verifying the refinement in linear time (see Section 3.3)

We call *fair transition* a transition of a set $F_i \in \mathcal{F}$.

Definition 2.9 (Fair execution (computation)) *An execution of $FTS = \langle TS, \mathcal{F}\rangle$ (also called a* computation *in [MP95]), obeys the fairness requirements $\mathcal{F}$, which means that: $(i \in [1..m] \ \wedge\ s \in Inf_s(\sigma) \ \wedge\ s \in In(F_i)) \Rightarrow \exists (s',\ a,\ s'') \cdot (s' \xrightarrow{a} s'' \ \in F_i \ \wedge\ s' \xrightarrow{a} s'' \in Inf_t(\sigma))$.*

The set Σ_{FTS} only contains computations. In other words, an execution σ is fair if it is true that "if some transition of any F_i is enabled infinitely often by σ, then some transition of F_i is taken infinitely often by the execution σ".

The executions of a transition system which do not satisfy fairness requirements are called *unfair executions.*

Property 2.10 *An execution σ that contains infinitely often two states s and s' which are in relation by a fair transition $s \xrightarrow{a} s'$ in F_i is a computation. So, $(s \in Inf_s(\sigma) \wedge s' \in Inf_s(\sigma) \wedge s \xrightarrow{a} s' \in F_i) \Rightarrow \exists(s_1, s_1') \cdot (s_1 \xrightarrow{a} s_1' \in F_i \wedge s_1 \xrightarrow{a} s_1' \in Inf_t(\sigma))$.*

Proof. Let σ be an execution that contains s and s' infinitely often. By the constraint (b) in Definition 2.8, σ contains a fair transition labelled by a infinitely often. Therefore, σ is a computation. □

Definition 2.11 (Fair Exiting Cycle) *Let $c = s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} \cdots \xrightarrow{a_{n-1}} s_0$ be a cycle of a fair transition system. The cycle c is a fair exiting cycle if there exists a transition $t = s_i \xrightarrow{a} s'$ of a fairness constraint F_j, such that $s' \neq s_i$, with $i \in [0 \cdots n]$ and $j \in [1 \cdots m]$.*

We call such a transition t an exit transition for a cycle c. By this definition we deduce that the computations of a system do not run around fair exiting cycles infinitely many times because they must take exit transitions.

2.3 Gluing Invariant

Let SP_V be a set of state propositions over V defined by the grammar

$p ::= ap \mid p \vee p \mid \neg p$ where $ap \in AP_V$.

Definition 2.12 (Validity of a state proposition) *A state proposition $p \in SP_V$ is valid*[1] *for a state s of a (fair) transition system (written as $s \models p$) if*

- $s \models ap$ *iff* $ap \in L(s)$,
- $s \models p_1 \vee p_2$ *iff* $s \models p_1$ *or* $s \models p_2$,
- $s \models \neg p$ *iff it is not true that* $s \vDash p$, *written as* $s \nvDash p$.

Let V_1 and V_2 be respectively the sets of variables of two transition systems TS_1 and TS_2. Let $SP_{V_{12}}$ be a set of state propositions over V_1 and V_2 defined by the following grammar:

$q ::= ap_1 \mid ap_2 \mid x_1 = x_2 \mid q \vee q \mid \neg q$ where $ap_1 \in AP_{V_1}$, $ap_2 \in AP_{V_2}$, $x_1 \in V_1$ *and* $x_2 \in V_2$.

Definition 2.13 (Validity of a state proposition on a pair of states) *A proposition $q \in SP_{V_{12}}$ is valid for a pair of states (written as $(s_1, s_2) \models q$) if*

- $(s_1, s_2) \models ap_1$ *iff* $ap_1 \in L_1(s_1)$,
- $(s_1, s_2) \models ap_2$ *iff* $ap_2 \in L_2(s_2)$,
- $(s_1, s_2) \models x_1 = x_2$ *iff*
  $\exists v \cdot (v \in D_{x_1} \ \wedge \ v \in D_{x_2} \ \wedge \ (x_1 = v) \in L_1(s_1) \ \wedge \ (x_2 = v) \in L_2(s_2))$,
- $(s_1, s_2) \models q_1 \vee q_2$ *iff* $(s_1, s_2) \models q_1$ *or* $(s_1, s_2) \models q_2$,
- $(s_1, s_2) \models \neg q$ *iff it is not true that* $(s_1, s_2) \models q$.

[1] we also say that "p holds on s"

We call gluing invariant, a state proposition of $SP_{V_{12}}$ which expresses how the variables from the abstract and the refined transition systems are linked together.

Definition 2.14 (Gluing invariant) *A state proposition $I_{12} \in SP_{V_{12}}$ is a* gluing *invariant of two (fair) transition systems $(F)TS_1$ over V_1 and $(F)TS_2$ over V_2, if I_{12} is an invariant on $S_1 \times S_2$ (i.e. $(s_1, s_2) \models I_{12}$ for all pairs (s_1, s_2) of $S_1 \times S_2$), and I_{12} is a total function from S_2 to S_1.*

We require I_{12} to be a total function from S_2 to S_1 because it allows partitioning TS_2 (see Section 3.1).

2.4 PLTL

Here, we define all future PLTL formulas with the two temporal operators, Next ($\bigcirc$) and Until ($\mathcal{U}$). We also use the following notations: $\Diamond P$ (eventually P) defined as $true\ \mathcal{U}\ P$, $\Box P$ (always P) defined as $\neg\Diamond\neg P$, and $P_1 \Rightarrow P_2$ defined as $\neg P_1 \vee P_2$. The composition of the temporal operators $\Box\Diamond$ means *infinitely often*.

In order to verify PLTL formulas on a (fair) transition system $(F)TS$, we extend Definition 2.12 to the PLTL semantics on the executions (or *computations*) in a standard way.

Definition 2.15 (PLTL) *Given PLTL formulas P, P_1, P_2 and an execution σ (or computation), we define P to be valid at the state s_j, $j \geq 0$, on an execution $\sigma = s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} \cdots, s_j \xrightarrow{a_j} \cdots$ (written as $(\sigma, j) \models P$) as follows:*

- $(\sigma, j) \models ap$ *iff* $ap \in L(s_j)$,
- $(\sigma, j) \models \neg P$ *iff it is not true that* $((\sigma, j) \models P)$, *written as* $(\sigma, j) \not\models P$,
- $(\sigma, j) \models P_1 \vee P_2$ *iff* $(\sigma, j) \models P_1$ *or* $(\sigma, j) \models P_2$,
- $(\sigma, j) \models \bigcirc P$ *iff* $(\sigma, j+1) \models P$,
- $(\sigma, j) \models P_1 \mathcal{U} P_2$ *iff* $\exists k \cdot (k \geq j \wedge (\sigma, k) \models P_2 \wedge\ \forall i \cdot (j \leq i < k \Rightarrow (\sigma, i) \models P_1))$.

When $(\sigma, 0) \models P$ we say that "P holds on σ" and we note $\sigma \models P$.

Now, we extend Definition 2.15 to transition systems.

Definition 2.16 (Validity of a PLTL formula on a (fair) transition system) *A PLTL formula P is valid for a (fair) transition system $(F)TS$, written as $(F)TS \models P$, if $\forall \sigma \cdot (\sigma \in \Sigma_{(F)TS} \Rightarrow \sigma \models P)$.*

The PLTL allows expressing many dynamics properties such as safety, liveness and fairness. Fairness assumptions expressed by the set $\mathcal{F}$ are described by the PLTL formula

$$f \stackrel{\text{def}}{=} \bigwedge_{i=1}^{m} \Box(\Box\Diamond \bigvee_{s \xrightarrow{a} s' \in F_i} \Big(\bigwedge_{ap \in L(s)} ap \Big) \Rightarrow \Diamond \bigvee_{s \xrightarrow{a} s' \in F_i} \Big(\bigwedge_{ap' \in L(s')} ap' \Big)) \tag{1}$$

which means that if a transition $t = s \xrightarrow{a} s'$ of a fairness constraint F_i is infinitely often enabled, then a transition of F_i must be taken infinitely often. This description is correct because of Property 2.10.

It is important to notice that the verification of a PLTL formula is the same on a transition system TS and on a τ-transition system τ-TS.

Property 2.17 *τ-$TS \models P$ iff $TS \models P$.*

Proof. The sequences of states which compose the executions in $\Sigma_{\tau-TS}$ and Σ_{TS} are the same by Definition 2.7. Since the PLTL satisfaction in Definition 2.15 is verified only on the sequences of states of the executions, Property 2.17 holds. □

A property of a system expressed as a PLTL formula is referred to as a *PLTL property*.

2.5 Büchi Automata and Executions Acceptance

We associate to a PLTL formula P a Büchi automaton denoted $\mathcal{B}_P$. The automaton $\mathcal{B}_P$ accepts all the executions on which P holds.

Definition 2.18 (Büchi automaton) *A Büchi automaton is a 5-tuple $\mathcal{B} = \langle q_0, Q, SP_V, \rightarrow_{\mathcal{B}}, A\rangle$ where:*

- *Q is a finite set of states ($q_0 \in Q$ is the initial state),*
- *$\rightarrow_{\mathcal{B}}$ is a finite set of transitions labelled by elements of SP_V : $\rightarrow_{\mathcal{B}} \subseteq Q \times SP_V \times Q$,*
- *$A \subseteq Q$ is the set of* accepting states *of the automaton.*

Similarly to the notion of execution of a transition system, an infinite alternating sequence of states of $\mathcal{B}$ and state propositions in SP_V is called a *run* of $\mathcal{B}$. We denote by $\Sigma_{\mathcal{B}}$ the set of all the runs of $\mathcal{B}$. A run of $\mathcal{B}$ is *accepting* if at least one of the accepting states appear infinitely often in the run.

Definition 2.19 (Execution Acceptance) *An execution $\sigma = s_0 \stackrel{a_0}{\rightarrow} s_1 \cdots s_i \stackrel{a_i}{\rightarrow} s_{i+1} \cdots \in \Sigma_{TS}$ is accepted by a run $\pi = q_0 \stackrel{p_0}{\rightarrow} q_1 \cdots q_i \stackrel{p_i}{\rightarrow} q_{i+1} \cdots \in \Sigma_{\mathcal{B}}$ if*

i) π is a run of $\mathcal{B}$ on σ: $\forall i \cdot ((0 \leq i \wedge q_i \stackrel{p_i}{\rightarrow} q_{i+1} \in \rightarrow_B) \Rightarrow s_i \models p_i)$,

ii) the run is accepting: $Inf_s(\pi) \cap A \neq \emptyset$.

3 Refinement

In this section, we express the refinement semantics as a relation between fair transition systems because we want to verify PLTL properties under fairness assumptions during the development by the refinement process. We improve the refinement relation between transition systems defined in [BJK00] by the fairness preservation clause in order to obtain the fair refinement relation.

We define the refinement relation as a state and action simulation which allows us to exploit a partition of the refined transition system state space into parts. With such a partition we are able to deal with the model checking blow-up by verifying PLTL properties of the refined system in a partitioned way. Moreover the fairness constraints of the environment make some abstract system behaviors fair, and these fairness constraints are preserved during the refinement steps. So, we want to verify PLTL properties under fairness assumptions in a partitioned way at the refined level.

3.1 Gluing Relation and State Space Partition

In this section, we consider two fair transition systems $FTS_1 = \langle S_{0_1}, S_1, Act_1, \rightarrow_1, L_1, \mathcal{F}_1\rangle$ over V_1 and $FTS_2 = \langle S_{0_2}, S_2, Act_2, \rightarrow_2, L_2, \mathcal{F}_2\rangle$ over V_2 giving the operational semantics of a system at two levels of refinement. The relation between the variables V_1 and V_2 is defined using a gluing invariant I_{12}. Our goal is to verify that FTS_1 is *refined* by FTS_2 (written as $FTS_1 \sqsubseteq_f FTS_2$) according to I_{12}.

In our approach, the refinement main features are as follows.

1. The refinement introduces *new* actions, so $Act_1 \subseteq Act_2$.
2. The refinement renames variables, so $V_1 \cap V_2 = \emptyset$.
3. The refinement introduces a *gluing* invariant, I_{12}.

Before defining the refinement relation, we define the gluing relation μ as a total function from S_2 to S_1. The refinement relation is the gluing relation restricted by additional clauses.

3.1.1 Gluing Relation and State Space Partition

We define a binary relation $\mu \subseteq S_2 \times S_1$ allowing us to express the relation between the states of two transition systems TS_1 and TS_2.

Definition 3.1 (Glued states) *Let I_{12} be the gluing invariant. The state $s_2 \in S_2$ is* glued *to $s_1 \in S_1$ by I_{12}, written $s_2\ \mu\ s_1$, if $(s_1, s_2) \models I_{12}$.*

Since μ is a total function from S_2 to S_1, we can define an equivalence relation $\sim_\mu$ between states of the refined transition system.

Definition 3.2 (Equivalence class) *Consider a state $s_1 \in S_1$. $EC(s_1)$ is an equivalence class of $S_2/_{\sim_\mu}$ if, for every state $s_2 \in EC(s_1)$, we have $s_2\ \mu\ s_1$.*

3.2 Refinement Relation

In this section, we define the refinement of two fair transition systems as a particular kind of simulation and we view it as computations containment. For that, we restrict μ into a function ρ_f which relates a refined fair transition system to one of its abstractions.

This relation allows us to distinguish some elements of the state space partition that we call parts. This partition is used either to prove an invariant of a part or to verify state propositions of a PLTL formula which are verified on a part (see Section 4). In order to describe the refinement, we keep the transitions of FTS_2 labelled over Act_1 but the *new* ones (from $Act_2 \setminus Act_1$) introduced by the refinement are considered as *non observable* τ moves. These τ moves hide the transitions of the parts viewed as transition systems (see Fig. 1). Let $Act_{1_\tau} \stackrel{\text{def}}{=} Act_1 \cup \{\tau\}$. In the above parts, it is certainly not desirable that τ moves take control forever. Therefore, infinite τ-executions are forbidden.

Let S_{c2} be the set of the states of S_2 (see Fig. 1) which are targets of a transition labelled by an abstract action, and are glued with the states of S_1 sources of at least a transition in a fairness constraint F_i. For example, S_{c2} is the set $\{s_2, w\}$ in Fig. 1. Any cycle at the refined level refining an abstract fair exiting cycle contains a state of the set S_{c2}. Let σ be a computation which runs around such a cycle c infinitely many often – this means it reaches a state of S_{c2}. The computation σ must leave c infinitely often in order to preserve the fairness constraints of the abstract level.

$$S_{c2} = \{s_2 \mid \exists(s, a_1').(s \stackrel{a_1'}{\rightarrow} s_2 \in \rightarrow_2)\ \wedge\ \exists(s_1,\ s_1',\ a_1).(s_1 \stackrel{a_1}{\rightarrow} s_1' \in \textstyle\bigcup_{i=1}^{m_1} F_{1i})\ \wedge\ F_{1i} \in \mathcal{F}_1\ \wedge\ s_2 \mu s_1\}.$$

Let $T_1 : S_{c2} \rightarrow 2^{\rightarrow_1}$ be a total function. We denote by $T_1(s_2)$ the set of the abstract fair transitions, such that their source states are glued with the state s_2 by the relation μ. For example, $T_1(s_2)$ is the set $\{t_1\}$ and $T_1(w)$ is the set $\{t\}$ in Fig. 1. The computations in which the state s_2 appears infinitely often, execute infinitely often the transitions of $\rightarrow_2$ which refine the transitions of $T_1(s_2)$.

$$T_1(s_2) = \{t_1 \mid t_1 = s_1 \stackrel{a_1}{\rightarrow} s_1' \wedge t_1 \in \bigcup_{i=1}^{m_1} F_{1i}\ \wedge\ F_{1i} \in \mathcal{F}_1\ \wedge s_2\ \mu\ s_1\}.$$

Let τ-$FTS_2 = \langle S_{0_2}, S_2, Act_{1_\tau}, \rightarrow_{2_\tau}, L_2, \mathcal{F}_2\rangle$ be the fair transition system resulting from FTS_2 and Act_1. Let s_1 be a state of FTS_1 ($s_1 \in S_1$) and s_2 be a state of FTS_2. Let a be a name of action of FTS_1, $a \in Act_1$.

Definition 3.3 (ρ_f relation) *Let FTS_1 and τ-FTS_2 be respectively two fair transition systems provided with a gluing invariant I_{12}. The relation $\rho_f \subseteq S_2 \times S_1$ is defined from I_{12} as the greatest binary relation included into μ satisfying the following clauses:*

1. ***strict refinement of abstract transition***
   $(s_2\ \rho_f\ s_1 \wedge s_2 \stackrel{a}{\rightarrow}_{2_\tau} s_2') \Rightarrow \exists s_1' \cdot (s_1 \stackrel{a}{\rightarrow}_1 s_1'\ \wedge\ s_2'\ \rho_f\ s_1'),$

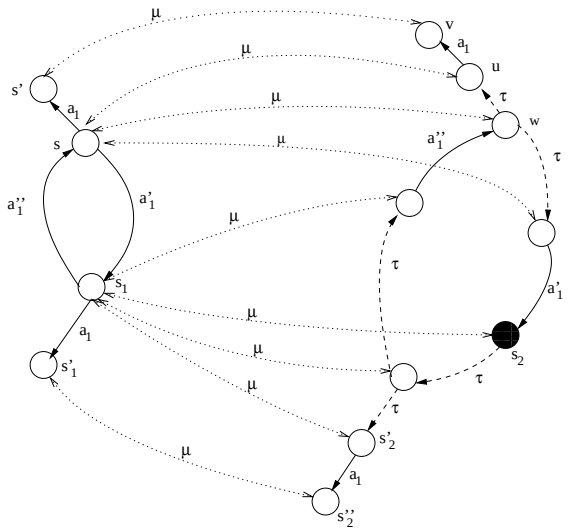

Figure 1: Refinement of a fair exiting cycle. The fair transitions at the abstract level are: $t = s \xrightarrow{a_1} s'$ and $t_1 = s_1 \xrightarrow{a_1} s'_1$, such that $F_{1i} = \{t, t_1\}$, $i \in [1 \dots m_1]$ and $\mathcal{F}_1 = \{F_{1i}\}$. The transitions t and t_1 are the exit transitions of the cycle at the abstract level. At the refined level the state s_2 belongs to the set S_{c2}. So, it is linked by the relation ρ_f to a state s_1 which is a source state of the fair transition t_1 ($t_1 \in F_{1i}$). Therefore $T_1(s_2) = \{t_1\}$.

2. ***τ-divergence freeness***

   $\forall \sigma \cdot (\sigma \in \Sigma_{\tau - FTS_2} \Rightarrow \forall m \cdot (m \in Act1_\tau^* \Rightarrow tr(\sigma) \neq m \cdot \tau^\omega))$,

3. ***stuttering of τ-transitions***

   $(s_2 \ \rho_f \ s_1 \ \wedge \ s_2 \xrightarrow{\tau}_{2\tau} s'_2) \Rightarrow s'_2 \ \rho_f \ s_1$,

4. ***abstract fairness preservation*** *(see Fig. 1)*

   $s_2 \ \rho_f \ s_1 \ \wedge \ s_2 \in S_{c2} \ \wedge \ s_1 \xrightarrow{a_1} s'_1 \in T_1(s_2) \ \wedge \ i \in [1 \dots m_1] \ \wedge \ s_1 \xrightarrow{a_1} s'_1 \in F_{1i} \ \wedge \ s_2 \in Inf_s(\sigma_2) \ \wedge \ \sigma_2 \in \Sigma_{FTS_2} \Rightarrow \exists (s, s', u, v).(s \xrightarrow{a_1} s' \in F_{1i} \ \wedge \ u \xrightarrow{a_1} v \in Inf_t(\sigma_2) \ \wedge \ u \ \rho_f \ s \ \wedge \ v \ \rho_f \ s')$,

5. ***non reduction of the abstract fairness constraints*** *(see Fig. 1)* $s_2 \ \rho_f \ s_1 \ \wedge \ s_2 \in S_{c2} \ \wedge \ s_1 \xrightarrow{a_1} s'_1 \in T_1(s_2) \Rightarrow \exists (\sigma_2, s'_2, s''_2).(\sigma_2 \in \Sigma_{FTS_2} \ \wedge \ s'_2 \xrightarrow{a_1} s''_2 \in \rightarrow_2 \ \wedge \ s'_2 \rho_f s_1 \ \wedge \ s''_2 \rho_f s'_1 \ \wedge \ (s_2 \in Inf_s(\sigma_2) \Rightarrow s'_2 \xrightarrow{a_1} s''_2 \in Inf_t(\sigma_2)))$.

Clauses 1-3 are already defined in [BJK00] in order to define the refinement relation between two transition systems. In this paper, we have added Clause 4 in order to define the refinement relation between fair transition systems. Notice that the presence of Clause 2 guarantees the monotonicity of an iterative construction of the relation ρ_f and, this way, the existence of this relation.

Clauses 1-3 of Definition 3.3 imply that all the τ-fragments (i.e. sequences of τ-transitions) are finite and are followed by a transition labelled in Act_1. Therefore, τ-livelocks are forbidden. In fair transition systems, the executions which run around τ-cycles infinitely many times without taking their exit transition infinitely often, are called τ-executions. New fairness assumptions must be introduced in the refined level in order to forbid τ-executions.

The aim of Clause 4 and Clause 5 is to preserve the abstract fairness constraints. At the abstract level, fairness constraints express restrictions in some executions of the system. The excluded executions are those

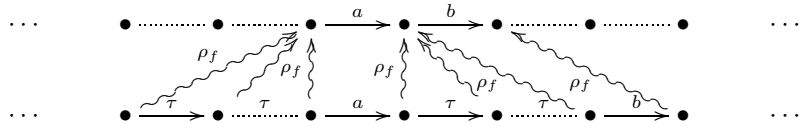

Figure 2: Included computations

which run around fair exiting cycles infinitely many times without taking their exit transitions infinitely often, they are not computations.

Clause 4 and Clause 5 are illustrated by Fig. 1 which represents the refinement of a fair exiting cycle. In this case, the relation ρ_f between τ-FTS_2 and FTS_1 satisfies Clause 4 if and only if every execution σ of τ-FTS_2 in which the state s_2 (or the state w)occurs infinitely often (s_2 belongs to a fair exiting cycle) activates infinitely often a transition of τ-FTS_2 which is simulated [2] by a transition in F_{1i} (because $t_1 \stackrel{\text{def}}{=} s_1 \stackrel{a_1}{\rightarrow} s'_1 \in F_{1i}$). Therefore, computations must activate infinitely often either the transition $s'_2 \stackrel{a_1}{\rightarrow} s''_2$ (see Fig. 1) (this transition is simulated by the transition t_1), or the transition $u \stackrel{a_1}{\rightarrow} v$ (this transition is simulated by the transition t and $t \stackrel{\text{def}}{=} s \stackrel{a_1}{\rightarrow} s' \in F_{1i}$).

The relation ρ_f between τ-FTS_2 and FTS_1 satisfies Clause 5 if and only if there exists executions of τ-FTS_2 which reaches infinitely often the states of S_{C2}, i.e. s_2 and w respectively (they are simulated by the states sources of the fair transitions t_1 and t respectively), and activates infinitely often the transitions of τ-FTS_2 which are simulated by t_1 and t respectively.

So, Clauses 4 and 5 show that any cycle c_2 of τ-FTS_2 which is simulated by a fair exiting cycle c_1, must satisfy the following conditions:

- c_2 must be a fair exiting cycle,
- if the exit transitions of c_1 are the transitions t and t_1 (see Fig. 1), then the computations which run around c_2 infinitely often, must activate infinitely often the transition $u \stackrel{a_1}{\rightarrow} v$ or the transition $s'_2 \stackrel{a_1}{\rightarrow} s''_2$ because they are simulated respectively by t and t_1,
- the cycle c_2 must have at least the same number of exit transitions as c_1.

Clause 4 and Clause 5 ensure that we have a refinement relation which preserves the abstract fairness constraints, and which also preserves the PLTL properties. It means that, a PLTL property satisfied at the abstract level is also verified at the refined level. For this, it is necessary that each computation of the refined level is simulated by a computation of the abstract system. This is guaranteed by Clauses 4 and 5. Abstract fairness assumptions are reformulated and introduced at the refined level in order to satisfy Clauses 4 and 5.

When the fair refinement relation holds between two fair transition systems FTS_1 and FTS_2, the executions of TS_2 which satisfy fairness assumptions also satisfy Clause 2 and Clause 4.

Remark 3.4 *Intuitively, the relation ρ_f implies computations containment. Given Σ_{FTS_1} and $\Sigma_{\tau-FTS_2}$, Clauses 1, 2 and 3 of Definition 3.3 mean that every execution of $\Sigma_{\tau-FTS_2}$ is linked to some execution in Σ_{FTS_1} (see Fig. 2). In other words, to every transition labelled a in FTS_1 corresponds a fragment σ' of an execution of FTS_2 where σ' is composed of a sequence of τ-transitions followed by a transition labelled by a.*

Definition 3.5 (Refinement) *A fair transition system $FTS_1 = \langle S_{0_1}, S_1, Act_1, \rightarrow_1, L_1, \mathcal{F}_1\rangle$ is refined by a fair transition system $FTS_2 = \langle S_{0_2}, S_2, Act_2, \rightarrow_2, L_2, \mathcal{F}_2\rangle$ provided with a gluing invariant I_{12}, written $FTS_1 \sqsubseteq_f FTS_2$, if $\forall s_2 \cdot (s_2 \in S_{0_2} \Rightarrow \exists s_1 \cdot (s_1 \in S_{0_1} \wedge s_2\ \rho_f\ s_1))$.*

Often initial states are designed ($S_{0_i} \subseteq S_i$). Then it is enough that ρ_f holds on the reachable state spaces of both systems. If there exist τ-cycles, they must be forbidden by the new fairness assumptions.

[2] we say that a transition $s_2 \stackrel{a}{\rightarrow} s'_2$ is simulated by a transition $s_1 \stackrel{a}{\rightarrow} s'_1$ when $s_2 \rho_f s_1$ and $s'_2 \rho_f s'_1$.

3.3 Analysis of the refinement verification

The algorithmic verification of refinement of finite transition systems can be effectively done by a joined exploration of the reachable state spaces. We have given a verification algorithm of the fair refinement in [CJ03, Cho03] (the proof of the algorithm was also given).

Let $\mid TS_1 \mid$ and $\mid TS_2 \mid$ be respectively the sizes of the transition system and one of its refinement. Suppose μ is a function. Verifying Clauses (1),(3) requires a parallel exploration of the systems TS_1 and TS_2. Same goes for Clauses (4),(5) because of the constraint (c) in Definition 2.8. Therefore, the complexity is $\mathcal{O}(|TS_1| + |TS_2|) \simeq \mathcal{O}(|TS_2|)$ (because generally $|TS_2| > |TS_1|$). For a finite TS_2, verifying Clause (2) requires a search for τ-cycles by exploring paths, of TS_2, and by following fair transitions. This verification does not change the complexity ($\mathcal{O}(|TS_2|)$).

However, if μ is not a function, the verification necessitates a joint enumeration of both systems. The complexity becomes in $\mathcal{O}(|TS_2| \times |TS_1|)$.

In the next section, we give an overview of how to verify PLTL properties without considering environment fairness constraints, by a partitioned model checking, as it is presented in [MMJ00, JMM01, Mas01].

4 Partitioned Model Checking

In this Section, we present the main results of an out-of-core model checking technique that we have developed in order to face the state explosion problem for the PLTL model checking. This technique has been presented in [MMJ00, JMM01] for transition systems without fairness assumptions.

In order to perform model checking on large transition systems, the partitioned verification technique relies on a simple idea: why not split the transition system into several smaller pieces, and perform the verification on each piece separately? The pieces are called *parts*. Parts are transition systems as well. The initial transition system is called the *global* transition system.

In order to have every transition in one part, the parts are constructed by partitioning the transitions of the global transition system. Some states may belong to two distinct parts: they can be the target state of a transition t in one part, and the initial state of a transition t' in another part.

To perform a partitioned verification is to verify a property on each part separately, and to conclude that it is globally true when it is true on every part.

Section 4.1 defines what is a property verifiable by parts, and exhibits a class of such PLTL properties. Section 4.2 proposes a partitioning of a transition system based on the refinement.

4.1 Properties Verifiable by Parts

4.1.1 Definition of a Property Verifiable by Parts

Consider a transition system split into a set of parts (transition systems) according to a partition of its set of transitions. Actually, it is enough that the parts are obtained by an overlapping of the transitions. Some PLTL properties have the property to be globally true when they are true on every part. We call such properties *verifiable by parts*.

Definition 4.1 (Property verifiable by parts) *Let P be a PLTL property. Let TS be a transition system, and let $\mathbb{M}$ be a partitioning of TS. The property P is verifiable by parts on TS if*

$$\forall M \cdot (M \in \mathbb{M} \Rightarrow M \models P) \Rightarrow TS \models P. \tag{2}$$

Remark 4.2 *We simply say P* is verifiable by parts*, instead of P* is verifiable by parts on TS*.*

4.1.2 A Class of PLTL Properties Verifiable by parts

Before we perform a partitioned verification, we have to make sure that the properties that we want to verify are verifiable by parts. That is, we have to prove (2) on every property. Let P be a PLTL property. Notice

that to prove "if P is true on every part, then it is true on the global transition system" is equivalent to prove "if P is false on the global transition system, then it is false on one part at least". That is, to prove (2) is equivalent to prove (3):

$$\neg(TS \models P) \Rightarrow \exists M \cdot \exists \sigma \cdot (M \in \mathbb{M} \wedge \sigma \in \Sigma_M \wedge \sigma \models \neg P). \tag{3}$$

By using Büchi automata, we give a sufficient condition for when a PLTL property is verifiable by parts. Consider a property P, and an execution σ on which P does not hold. Suppose that there is a state s in σ such that every fragment of σ starting in s violates P. Then the part containing s violates P since it contains a fragment of σ starting in s. The same idea works when a unique transition is the cause of such a violation of the property. Because every state (and every transition) necessarily belongs to one part, we know that the property does not hold on the part that contains this state (or transition).

We define a class $\mathcal{B}_{mod}$ of Büchi automata, and we prove that every PLTL property whose negation defines a language that is recognized by an automaton in $\mathcal{B}_{mod}$ is a property verifiable by parts. An automaton in $\mathcal{B}_{mod}$ has every accepting state leading only to an accepting state, and there is at most one intermediate state between the initial state and any accepting state. Moreover, there is a loop labelled by `True` on the initial state.

Definition 4.3 (Class $\mathcal{B}_{mod}$ of Büchi automata) *Let $\mathcal{B} = \langle q_0, Q, SP_V, \rightarrow_{\mathcal{B}}, A\rangle$ be a Büchi automaton. We have $\mathcal{B} \in \mathcal{B}_{mod}$ if*

1. *there is a loop labelled `True` on the initial state: $q_0 \xrightarrow{\texttt{True}} q_0 \in \rightarrow_{\mathcal{B}}$,*
2. *for any run $\pi = \ q_0 \xrightarrow{p_0} q_1 \xrightarrow{p_1} q_2 \dots q_i \xrightarrow{p_i} q_{i+1} \dots \ \in \Sigma_{\mathcal{B}}$*

$$\exists k \cdot (k > 0 \wedge \forall i \cdot ((0 \leq i < k \Rightarrow q_i = q_0) \wedge (i > k \Rightarrow q_i \in A))), \tag{4}$$

3. *$\forall (q, p, q') \cdot (q \xrightarrow{p} q' \in \rightarrow_{\mathcal{B}} \wedge q' \in A \Rightarrow \exists (p', q'') \cdot (q' \xrightarrow{p'} q'' \in \rightarrow_{\mathcal{B}} \wedge p \Rightarrow p'))$.*

Theorem 4.4 *All the PLTL properties whose negation defines a language that is recognized by a Büchi automaton in the class $\mathcal{B}_{mod}$ are properties verifiable by parts.*

Proof. Let P be a PLTL property and let TS be a transition system on which P does not hold. We have $\neg(TS \models P)$, that is:

$$\exists \sigma \cdot (\sigma \in \Sigma_{TS} \wedge \sigma \models \neg P).$$

Let $\mathcal{B}_{\neg P} \in \mathcal{B}_{mod}$ be a Büchi automaton that recognizes the language of $\neg P$. There exists a run $\pi = q_0 \xrightarrow{p_0} q_1 \xrightarrow{p_1} q_2 \dots q_i \xrightarrow{p_i} q_{i+1} \dots \in \Sigma_{\mathcal{B}_{\neg P}}$ of $\mathcal{B}_{\neg P}$ on σ on which (4) holds. With an index k as defined in Formula (4), we consider s_{k-1} and s_k, respectively the $(k-1)$-th and k-th states in the execution σ. With any partitioning of TS, the transition $s_{k-1} \xrightarrow{a_{k-1}} s_k$ necessarily belongs to a part M, and the state s_{k-1} is reachable with transitions of M from an initial state s'_0 of M. Let

$$\sigma' = s'_0 \xrightarrow{a'_0} s'_1 \cdots s_{k-1} \xrightarrow{a_{k-1}} s_k \cdots$$

be a fragment of an execution of TS such that the suffix $s_{k-1} \xrightarrow{a_{k-1}} s_k \cdots$ is common to σ and σ', and such that all the transitions appearing in the prefix $s'_0 \xrightarrow{a'_0} s'_1 \cdots s_{k-1} \xrightarrow{a_{k-1}} s_k$ of σ' are transitions of M.

Consider the run

$$\pi' = q_0 \xrightarrow{\texttt{True}} q_0 \cdots q_0 \xrightarrow{\texttt{True}} q_{k-1} \xrightarrow{p_{k-1}} q_k \xrightarrow{p_k} q_{k+1} \cdots$$

of $\mathcal{B}_{\neg P}$ where the suffix $q_{k-1} \xrightarrow{p_{k-1}} q_k \cdots$ is common to π and π'. Such a run exists as $\mathcal{B}_{\neg P} \in \mathcal{B}_{mod}$ and as $q_{k-1} = q_0$ by construction. Moreover, the run π' is accepting on σ' because it "stays" on q_0 until it accepts the suffix $s_{k-1} \xrightarrow{a_{k-1}} s_k \cdots$ the same way it accepted it in σ.

There are two cases.

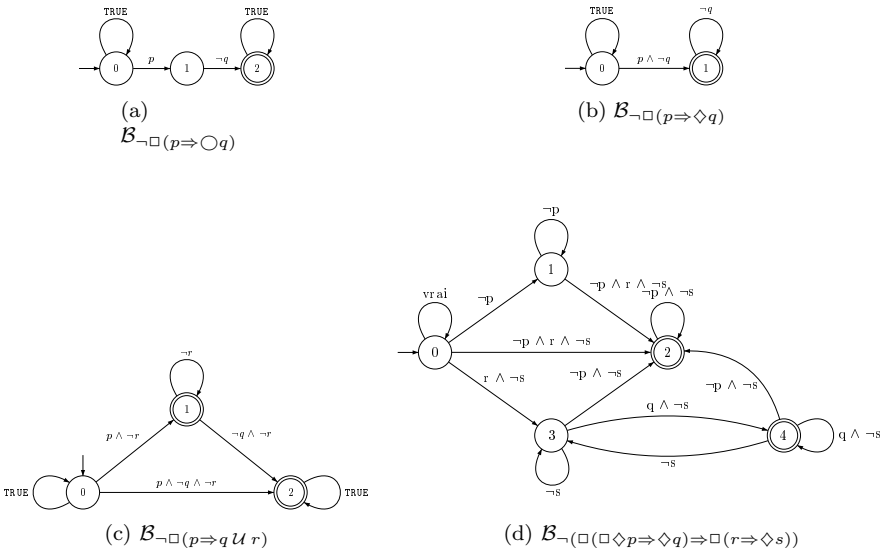

Figure 3: Some Büchi automata

1. σ' is an execution of M. Then σ' is accepted by $\mathcal{B}_{\neg P}$ and so $\neg(M \models P)$.

2. σ' is not an execution of M: there is a state s_c $(c \geq k)$ such that the prefix $s'_0 \stackrel{a'_0}{\rightarrow} s'_1 \cdots s_{k-1} \stackrel{a_{k-1}}{\rightarrow} s_k \cdots s_{c-1} \stackrel{a_{c-1}}{\rightarrow} s_c$ of σ' is a fragment of an execution

$$\sigma'' = s'_0 \stackrel{a'_0}{\rightarrow} s'_1 \cdots s_{k-1} \stackrel{a_{k-1}}{\rightarrow} s_k \cdots s_{c-1} \stackrel{a_{c-1}}{\rightarrow} s_c \stackrel{\texttt{Skip}}{\rightarrow} s_c \stackrel{\texttt{Skip}}{\rightarrow} s_c \cdots$$

of M. The transition $s_c \stackrel{a_c}{\rightarrow} s_{c+1}$ is not in M.

Consider

$$\pi'' = q_0 \stackrel{\texttt{True}}{\rightarrow} q_0 \cdots q_c \stackrel{p_c}{\rightarrow} q_{c+1} \stackrel{p'_0}{\rightarrow} q'_1 \stackrel{p'_1}{\rightarrow} q'_2 \cdots$$

an accepting run of $\mathcal{B}_{\neg P}$ where the prefix $q_0 \cdots q_{c+1}$ is common to π' and π''. As $\mathcal{B}_{\neg P} \in \mathcal{B}_{mod}$, then $\forall j \cdot (j \geq 1 \Rightarrow q'_j \in A)$. Moreover, from Clause 3 in Definition 4.3, $p_c \Rightarrow p'_0$ and $\forall j \cdot (j \geq 1 \Rightarrow (p'_{j-1} \Rightarrow p'_j))$. As a consequence, $\forall j \cdot (j \geq 0 \Rightarrow (p_c \Rightarrow p'_j))$. Thus, π'' is an accepting run on σ'', i.e. $\forall i \cdot (i \geq k-1 \wedge i \leq c \Rightarrow s_i \models p_i) \wedge \forall j \cdot (j \geq 1 \Rightarrow s_c \models p'_j)$. As the execution σ'' is accepted by $\mathcal{B}_{\neg P}$, then $\neg(M \models P)$.

□

The class $\mathcal{B}_{mod}$ contains some safety properties such as $\Box p$ or $\Box(p \Rightarrow \bigcirc q)$ (see Fig. 3(a)), all reachability properties such as $\Box\neg p$, and some liveness properties such as $\Box(p \Rightarrow \Diamond q)$ and $\Box(p \Rightarrow q\mathcal{U}\,r)$ (see Fig. 3(b) and Fig. 3(c)). However it does not contain liveness properties under fairness assumptions such as $\Box(\Box\Diamond p \Rightarrow \Diamond q) \Rightarrow \Box(r \Rightarrow \Diamond s)$ (see Fig. 3(d)).

4.2 A partitioning based on refinement

Consider a property P that is verifiable by parts, and a transition system TS on which P globally holds. Actually, it is very likely that with an inappropriate partitioning $\mathbb{M}$ of TS, there is a part M of $\mathbb{M}$ on which P does not hold. In this case, we can not conclude that P is globally true from the partitioned verification. In other words, the fact that a property globally true is also true on every part depends on the way the global transition system is partitioned.

As an heuristic to the choice of an appropriate partitioning, we propose to partition a transition system according to the refinement process.

At every step of the refinement, the specifier introduces new actions and refines the old ones. We propose that the specifier also introduces the PLTL properties that can be observed thanks to these new actions at this very level of refinement. The properties verified at the former levels of refinement need not to be verified again since PLTL properties are preserved by refinement [DJK03].

The new properties are likely to be observed on the "successions" of new actions, that is on the τ-executions. Intuitively, a property that could be observed on a succession of more than one old action is not a new property.

According to this idea, we propose that the parts contain the τ-executions of a refined system. Parts constructed in this way are called *refinement based parts*. We have proposed in [JMM01] a definition of such parts. The parts are built as follows (see Remark 3.4): to every state of the abstract system corresponds a part in the refined system. Let s_1 be a state of the abstract system. The part corresponding to s_1 is made of all the transitions of the refined system that have a state of $EC(s_1)$ as a source state. The target state s' of such a transition is

- either a state itself glued to s_1 – then the transition represents the occurrence of a new event, or of an old event if it refines an abstract transition $s_1 \xrightarrow{a} s_1$,
- or a state not glued to s_1 – then the transition represents the occurrence of an old event, and means that the "end" of the module is reached: s' has no successor in the part other than itself by the virtual transition labelled `Skip` (see Remark 2.2).

The definition that we propose in this paper is slightly different because we are in the context of the refinement of fair transition systems. Thus, Definition 4.5 defines a part of a refined fair transition system according to the fair refinement relation ρ_f. A part is a transition system.

Let us first give the intuition of what are the initial states, the states and the transitions of such a part. Consider an abstract fair transition system FTS_1 and a fair transition system FTS_2 that refines FTS_1. Let s_1 be a state of FTS_1. We define the part TS_M corresponding to s_1. We define Y as being the greatest set of states of FTS_2 that are successors of a state in $EC(s_1)$ by taking a transition labelled with an abstract action: $Y = \{\ s' \mid s_2 \xrightarrow{a} s' \in \rightarrow_2 \wedge\ s_2 \in EC(s_1) \wedge s' \notin EC(s_1)\}$. We define $FS(Y)$ as being the greatest set of states of FTS_2 that are reachable from a state of Y by taking only fair transitions: $FS(Y) = \{s_j \mid 1 \leq j \leq n \wedge \exists \sigma \cdot (\sigma \in \Sigma_{FTS_2} \wedge (s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} \ldots \xrightarrow{a_{i-1}} s_i \ldots \xrightarrow{a_{n-1}} s_n) \subset \sigma \wedge s_0 \in Y \wedge \forall i \cdot (s_i \xrightarrow{a_i} s_{i+1} \in \bigcup_{j=1}^{m_2} F_{2j}))\}$.

The initial states of TS_M are the states of $EC(s_1)$ that are either initial states of FTS_2, or target states of a transition labelled with an abstract action whose source state is not in $EC(s_1)$.

The transitions of TS_M are:

- the transitions of FTS_2 that have a state of $EC(s_1)$ as a source state,
- the fair transitions of FTS_2 that have a state of $Y \cup FS(Y)$ as a source state,
- the transitions labelled with `Skip` that are added as loops on the states of TS_M which have no successors among the states of TS_M.

Definition 4.5 (Refinement Based Part of a Fair Transition System) *Let $FTS_1 = \langle S_{0_1}, S_1, Act_1, \rightarrow_1, L_1, \mathcal{F}_1\rangle$ be a fair transition system which is refined by a fair transition system $FTS_2 = \langle S_{0_2}, S_2, Act_2, \rightarrow_2, L_2, \mathcal{F}\rangle$ ($FTS_1 \sqsubseteq_f FTS_2$). Consider $s_1 \in S_1$ and $EC(s_1)$, an equivalence class of $S_2/_{\sim_\mu}$. The part based on $EC(s_1)$ is a transition system $TS_M = \langle S_{0_M}, S_M, Act_M, \rightarrow_M, L_M\rangle$ such that:*

- $S_{0_M} = \{s_2 \in EC(s_1) \mid s_2 \in\ S_{0_2} \vee \exists s \cdot \exists a \cdot (s \xrightarrow{a} s_2 \in \rightarrow_2 \wedge\ s \notin EC(s_1))\}$.
- $S_M = \{s_2 \in EC(s_1)\} \cup Y \cup FS(Y)$.
- $\rightarrow_M = \{s_2 \xrightarrow{a} s' \in \rightarrow_2 \mid s_2 \in EC(s_1)\}\ \cup\ \{s \xrightarrow{a} s' \in \rightarrow_2 \mid s \in (Y\ \cup\ FS(Y))\ \wedge\ s' \in FS(Y)\}\ \cup\ \{s \xrightarrow{Skip} s \mid s \in S_M\ \wedge\ \forall (s', a) \cdot (s \xrightarrow{a} s' \in \rightarrow_2 \Rightarrow s' \notin S_M)\}$.

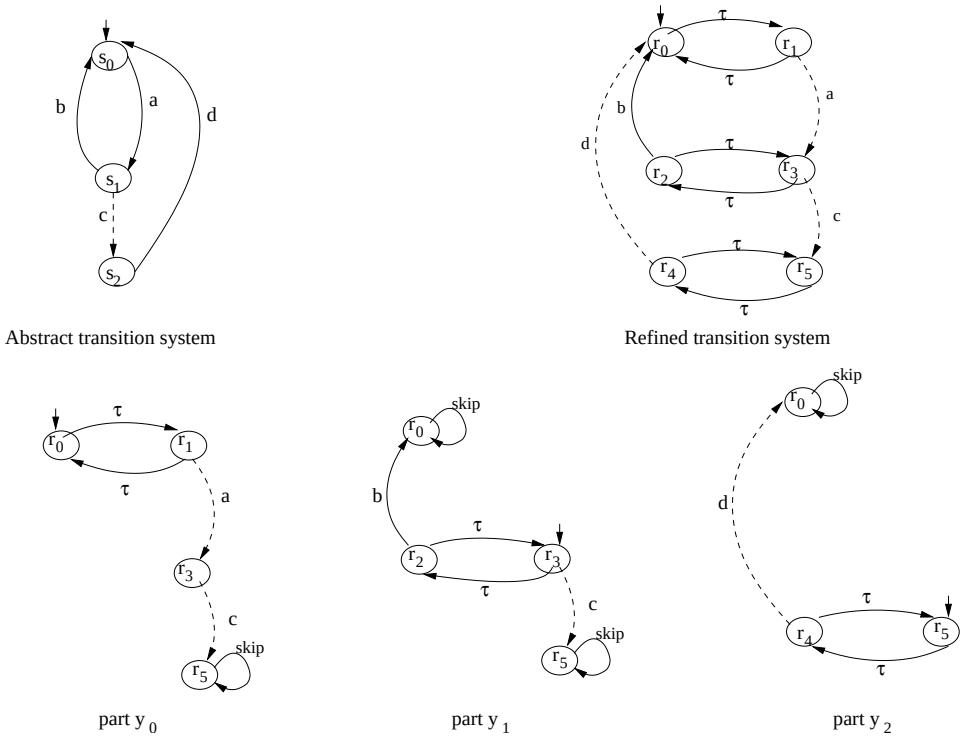

Figure 4: Example of a refinement based parts

- *Act_M is the restriction of Act_2 to the labels of $\rightarrow_M$, augmented with* `Skip`*.*
- *L_M is the restriction of L_2 on the states of S_M.*

Notice that there are at most as many parts in the refined transition system as states in the abstract transition system. As a consequence, the number of states of the abstract model can be used as a parameter for the user to control either the number, or the size of the parts to be model-checked.

Remark 4.6 *The properties that are not verifiable by parts for all possible partitionings, but only relatively to the partitioning as presented in Definition 4.5, are called* verifiable by refinement based parts.

Example of refinement based partition according to Definition 4.5

In Fig. 4, we present an example of three parts y_0, y_1, and y_2 obtained according to Definition 4.5 from the refined transition system. We suppose that the states of the refined system and the abstract system are glued as : r_0 and r_1 with s_0, r_2 and r_3 with s_1, and r_4 and r_5 with s_2. The fair transitions are represented by dashed arrows in Fig. 4.

For example the part y_0 is the transition system composed of the states $EC(s_0) = \{r_0,\ r_1\}$, $Y = \{r_3\}$ and $FS(Y) = \{r_5\}$. It has one initial state r_0. It contains all τ-transitions between the states of $EC(s_0)$, the exit transition $r_1 \xrightarrow{a} r_3$, and the fair transition $r_3 \xrightarrow{c} r_5$. As this transition system deadlocks in r_5, we extend it with the transition $r_5 \xrightarrow{Skip} r_5$.

Next section describes how the partitioned verification applies for the verification of PLTL properties under fairness assumptions.

5 Partitioned Model Checking under Fairness Assumptions

This section contains the main contribution of the paper. We study the verification by the partitioned model checking of PLTL properties when the description of the environment uses fairness constraints. We show that partitioned model checking on transition systems can be used under fairness assumptions.

We do not include the fairness of the environment in the transition system, but it is integrated as an assumption of the property to verify.

The verification by model checking of a PLTL property P under fairness assumptions expressed by a PLTL formula f, on a transition system TS, supposes verifying by model checking the property $f \Rightarrow P$ on TS. The question "does TS provided with fairness assumptions f satisfies P ?" is written as: $TS \models f \Rightarrow P$?

The Büchi automaton of $\neg(f \Rightarrow P)$ does not satisfy the sufficient condition allowing its verification in a partitioned way (see Section 4 and Fig. 3d). However, we prove that the verification of the formula $f \Rightarrow P$ becomes verifiable by refinement based parts (see Definition 4.5) , i.e. $\forall M.(M$ *is a refinement based part* $\Rightarrow M \models f \Rightarrow P) \Rightarrow TS \models f \Rightarrow P$ when P is verifiable by parts. This is equivalent to say that if, for all M such that M is a refinement based part, each computation of M satisfies P then all the computations of TS satisfies P. Indeed, since the unfair executions of M and TS satisfy the formula $f \Rightarrow P$ because they do not satisfy f, we deduce that the verification of $f \Rightarrow P$ on M (refinement based part) and on TS requires only the verification of P on the computations (fair executions) of M and TS.

Now, let us prove that each computation of TS is a concatenation of fragments which are all prefix of computations in refinement based parts.

Lemma 5.1 *Suppose $FTS_1 \sqsubseteq_f FTS_2$. Each computation σ of FTS_2 can be decomposed into fragments which are prefix of computations in refinement based parts of FTS_2. More precisely, such a computation is, either a suffix of σ, or is a fragment of σ which ends by a finite sequence of fair transitions followed by an infinite sequence of skip transitions.*

Proof. As the fair refinement is a τ-simulation of τ-FTS_2 by FTS_1, each computation σ of τ-FTS_2 is such that

$$\sigma = \ s_0 \stackrel{\tau*}{\rightarrow} s_{i_1-1} \stackrel{a_{i_1}}{\rightarrow} s_{i_1} \stackrel{\tau*}{\rightarrow} s_{i_2-1} \stackrel{a_{i_2}}{\rightarrow} s_{i_2} \stackrel{\tau*}{\rightarrow} s_{i_3-1} \stackrel{a_{i_3}}{\rightarrow} s_{i_3} \ldots$$

where a_{i_j} are abstract actions. The decomposition by Definition 4.5 is such that each finite fragment of σ, $\varphi_f = \ s_{i_{j-1}} \stackrel{\tau*}{\rightarrow} s_{i_j-1} \stackrel{a_{i_j}}{\rightarrow} s_{i_j}$ is a prefix of a computation of a refinement based part such as, either

- $\sigma' = \ s_{i_{j-1}} \stackrel{\tau*}{\rightarrow} s_{i_j-1} \stackrel{a_{i_j}}{\rightarrow} s_{i_j} \stackrel{c*}{\rightarrow} s_c \stackrel{\tau*}{\rightarrow} s_{i_j-1} \stackrel{a_{i_j}}{\rightarrow} s_{i_j} \stackrel{c*}{\rightarrow} s_c \ldots$, where $c \in Act_1 \cup \{\tau\}$ and $s_{i_j} \stackrel{c*}{\rightarrow} s_c$ is a finite sequence of fair transitions and states s_c are not source states of a fair transition -here σ' refines a suffix of a computation of a global system in the abstract level which run around a cycle infinitely many times -, or
- $\sigma'' = \ s_{i_{j-1}} \stackrel{\tau*}{\rightarrow} s_{i_j-1} \stackrel{a_{i_j}}{\rightarrow} s_{i_j} \stackrel{c*}{\rightarrow} s_c \stackrel{Skip}{\rightarrow} s_c \stackrel{Skip}{\rightarrow} s_c \ldots$.

Obviously σ' is a computation since σ is a computation and σ' is a suffix of σ. Also σ'' is a computation since, by construction, it is a fragment of the computation σ prolongated by a finite sequence of fair transitions followed by an infinite sequence of *skip* transitions. □

Correctness of the Partitioned Model Checking under Fairness Assumptions

In this section, we show that model checking by parts under fairness assumptions is sound. For that it is necessary that the parts are obtained according to Definition 4.5.

Theorem 5.2 *Let $FTS = \langle TS, \mathcal{F} \rangle$ be a fair transition system which refines an abstract fair transition system. Let $\mathbb{M}$ be the set of refinement based parts accordingly to Definition 4.5. Let f be the PLTL formula which expresses the fairness assumptions of FTS. If P is a PLTL formula such that $\mathcal{B}_{\neg P} \in \mathcal{B}_{mod}$ (as such, it is verifiable by parts on TS), then the property $f \Rightarrow P$ is verifiable by refinement based parts on TS.*

Proof. Recall that in order to prove that the formula $f \Rightarrow P$ is verifiable by refinement based parts on TS, we must prove the following: if, for all M such that M is refinement based part, each computation of M satisfies P then all the computations of TS satisfies P.

Let σ be a computation of TS. By Lemma 5.1, each fragment of a computation σ prefixes a computation like σ' or σ'' of a refinement based part. Therefore all the computations of a refinement based parts are extensions of all the fragments of the computations belonging to a global system.

In the proof of Theorem 4.4 in Section 4.1.2 it is shown that when a property P belong to the class of properties verifiable by parts, and if the extensions of all the fragments of an execution σ_i (in a global system) satisfy P, then P holds on σ_i.

So, when each computation of a refinement based part satisfies P and since P is verifiable by parts, we conclude by Theorem 4.4 that P holds on σ. Which means that $f \Rightarrow P$ is verifiable by refinement based parts on TS. □

6 Example of the Protocol T=1

In this section, we present the example of the protocol T=1 [edNE92] in order to illustrate how to verify PLTL properties under fairness assumptions in a partitioned way. We also defined how to express the fairness of an environment in a B event system. We give the B event systems enriched with fairness assumptions, of the protocol at the abstract and at the refined level. We also give the fair transition system which is the semantics of the abstract event system and the fair transition system which expresses the semantics at the refined level.

We also use this example to show that without using fair refinement to obtain the set of parts, the PLTL properties under fairness assumptions are not verifiable by parts.

6.1 Abstract Specification under Fairness Assumptions

Figure 5 represents the abstract B event system of a half duplex communication protocol between a chip integrated card and a card reader. At this level of specification, we consider only the alternation of exchange of messages between the chip card and the reader.

Figure 6 represents the abstract transition system of the protocol. In this figure, each state is decorated with the value of the state variables. $Cstatus_1$ indicates if the chip card is inserted or not in the reader. $Sender_1$ indicates the device which will send the next message, the chip card or the reader. The character '?' or '!' in the reader indicates respectively that the reader is the receiver or the sender device. The state labelling function, called L_1 is the following: $L_1 = \{$

$s_0 \mapsto \{Cstatus_1 = in, Sender_1 = reader\}$,

$s_1 \mapsto \{Cstatus_1 = in, Sender_1 = card\}$,

$s_2 \mapsto \{Cstatus_1 = out, Sender_1 = card\}$,

$s_3 \mapsto \{Cstatus_1 = out, Sender_1 = reader\}\}$.

At the initial state, the chip card is inserted in the reader, and sending a message must be done by the reader. The protocol evolves by the action of four events:

- *Rsends:* the reader sends a message,
- *Csends:* the chip card sends a message,
- *Eject:* the chip card is ejected,
- *Cinsert:* the chip card is inserted.

MACHINE teg1

SETS

$SENDER = \{card, reader\}$; $CARD - STATE = \{in, out\}$

VARIABLES

$Sender_1$, $Cstatus_1$

INVARIANT

$Sender_1 \in SENDER \wedge Cstatus_1 \in CARD - STATE$

INITIALISATION

$Sender_1 := reader \parallel Cstatus_1 := in$

EVENTS

Rsends $\hat{=}$ **SELECT** $Sender_1 = reader \wedge Cstatus_1 = in$
THEN $Sender_1 := card$**END**;

Csends $\hat{=}$ **SELECT** $Sender_1 = card \wedge Cstatus_1 = in$
THEN $Sender_1 := reader$ **END**;

Eject $\hat{=}$ **SELECT** $Cstatus_1 = in$
THEN $Cstatus_1 := out$ **END**;

Cinsert $\hat{=}$ **SELECT** $Cstatus_1 = out$
THEN $Sender_1 := reader \parallel Cstatus_1 := in$ **END**;

FAIRNESS= {Eject };

END

Figure 5: Abstract B specification of the protocol T=1 under fairness assumptions

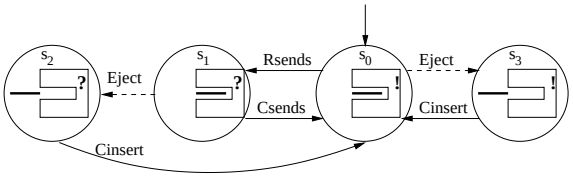

Figure 6: Abstract fair transition system of the protocol T=1

We assume that all the applications which use the protocol do not request the transport of an infinite sequences of messages. This is a fairness constraint which comes from the environment. It ensures that the transmission of the messages between the card and the reader terminates by the ejection of the card.

In B event systems we proposed the following extension. The fairness assumptions are written (e *if* p) [BCJK01], where e is the name of an event and p is a predicate characterizing the states in which e cannot be avoided when it is enabled infinitely often. When an event must always be fair, the fairness assumption is only the event name.

The fairness assumption on the environment protocol is defined by the clause *FAIRNESS*={*Eject* } (see Fig. 5). In the abstract fair transition system of the protocol, the fairness assumption is expressed by the set $\mathcal{F}_1 = \{F_{11}\}$ where $F_{11} = \{s_0 \stackrel{Eject}{\rightarrow} s_3,\ s_1 \stackrel{Eject}{\rightarrow} s_2\}$.

The set of states $\{s_0, s_1\}$ is defined by the following expression : $Cstatus_1 = in$. The expression $Cstatus_1 = out$ defines the set of states $\{s_2, s_3\}$. So, the verification of a PLTL property P on the transition system must be done under this fairness assumption which is expressed by the PLTL formula:

$$\Box(\Box\Diamond\ (Cstatus_1 = in)\ \Rightarrow\ \Diamond\ (Cstatus_1 = out)).$$

This formula is a simplification of the formula obtained from the general formula 1 presented in Section 2 that would be:

$\Box(\Box\Diamond\ (Cstatus_1 = in\ \wedge\ Sender_1 = reader) \vee (Cstatus_1 = in\ \wedge\ Sender_1 = card))\ \Rightarrow\ \Diamond\ ((Cstatus_1 = out\ \wedge\ Sender_1 = reader) \vee (Cstatus_1 = out\ \wedge\ Sender_1 = card))$.

6.2 Refined Specification under Fairness Assumptions

Figure 7 describes the refined B event specification of the protocol. At this level of specification, we view the messages as a sequence of blocks ended by a last block (value *lb*). A block sent (value *bl*) is acknowledged by an acknowledgement block (value *ackb*). We call *frame* these three types of exchanged information.

After a last block is sent by one of the devices, the other device answers with a sequence of blocks ending by a last block unless the card is ejected. These exchanges of messages alternate until the card is ejected.

The fair transition system shown in Fig. 8 describes the refined behavior of the protocol. The states of the transition system are decorated by the value of variables $CardF_2$ and $ReaderF_2$ which describe the type of the last frame sent respectively by the chip card and the reader, $SenderF_2$ which describes the device that will send the next frame, and $Cstatus_2$ which does the same thing as the variable $Cstatus_1$.

The gluing invariant is a part of the invariant in Fig. 7:
$I_{12} = ((Cstatus_2 = Cstatus_1)\ \wedge$
$((ReaderF_2 = bl\ \vee\ ((CardF_2 = ackb\ \vee\ CardF_2 = lb)\ \wedge\ SenderF_2 = reader))\ \Leftrightarrow\ (Sender_1 = reader))\ \wedge$
$((CardF_2 = bl\ \vee\ ((ReaderF_2 = ackb\ \vee\ ReaderF_2 = lb)\ \wedge\ SenderF_2 = card))\ \Leftrightarrow\ (Sender_1 = card)))$.

With this gluing invariant, the states of the two transition systems in Fig. 6 and Fig. 8 are glued as follows:

- $r_0, r_2, r_3, r_4, r_{10}$ are glued with the state s_0,
- $r_5, r_7, r_8, r_9, r_{12}$ are glued with the state s_1,
- r_6, r_{13} are glued with the state s_2,
- r_1, r_{11} are glued with the state s_3.

From these four equivalence classes (see Section. 4.2) we construct the four parts described in Fig. 9, Fig. 10, Fig. 11 and Fig. 12 (see the appendix).

The state labelling function, called L_2 is the following:$L_2 = \{$

$r_0 \mapsto \{Cstatus_2 = in,\ SenderF_2 = reader,\ ReaderF_2 = lb,\ CardF_2 = lb\}$,

$r_1 \mapsto \{Cstatus_2 = out,\ SenderF_2 = reader,\ ReaderF_2 = lb,\ CardF_2 = lb\}$,

$r_2 \mapsto \{Cstatus_2 = in,\ SenderF_2 = card,\ ReaderF_2 = bl,\ CardF_2 = lb\}$,

$\ldots\ \}$.

The old events *Rsends* and *Csends* terminate the emission of a message by sending the last block. We have reinforced the guard of the event *Eject* in order to forbid the ejection of the chip card during the transmission of a message. We have introduced four new events to take the transmission of blocks and acknowledgement into account:

- *Rblocksends:* the reader sends a block,
- *Cblocksends:* the card sends a block,
- *Racksends:* the reader sends an acknowledgement,
- *Cacksends:* the card sends an acknowledgement.

The fairness assumptions in the refined level are defined by the declaration set *FAIRNESS*={*Eject, Csends if* $(CardF_2 = bl)$, *Rsends if* $(ReaderF_2 = bl)$}. The fairness assumption $h_1 \stackrel{\text{def}}{=} Eject$ reformulates the abstract fairness assumption. It indicates that the end of the transmission is unavoidable. The fairness assumptions $h_2 \stackrel{\text{def}}{=} Csends\ if\ (CardF_2 = bl)$ and $h_3 \stackrel{\text{def}}{=} Rsends\ if\ (ReaderF_2 = bl)$ are new fairness assumptions which express that the messages sent by the chip card or the reader contain a finite number of blocks. They allow to go out of livelocks as shown in Fig. 8. These assumptions must be satisfied by all the possible environments of the protocol.

The verification of a PLTL property P on the transition system must be done under these fairness assumptions, expressed by the following PLTL formulae:

- $f_1' \stackrel{\text{def}}{=} \Box(\Box \Diamond (Cstatus_2 = in \wedge ((SenderF_2 = reader \wedge CardF_2 = lb) \vee (SenderF_2 = card \wedge ReaderF_2 = lb))) \Rightarrow \Diamond Cstatus_2 = out)$ expresses h_1,
- $f_{21} \stackrel{\text{def}}{=} \Box(\Box \Diamond(SenderF_2 = card \wedge CardF_2 = bl) \Rightarrow \Diamond CardF_2 = lb)$ expresses h_2,
- $f_{22} \stackrel{\text{def}}{=} \Box(\Box \Diamond (SenderF_2 = reader \wedge ReaderF_2 = \text{bl}) \Rightarrow \Diamond ReaderF_2 = lb)$ expresses h_3.

The fairness assumptions are defined by the sets of fair transitions (represented in Fig. 8 by the dashed arrows) $\mathcal{F}_2 = \{F_1', F_{21}, F_{22}\}$ where :

- $F_1' \stackrel{\text{def}}{=} \{r_0 \stackrel{Eject}{\rightarrow} r_1,\ r_{10} \stackrel{Eject}{\rightarrow} r_{11},\ r_{12} \stackrel{Eject}{\rightarrow} r_{13}, r_5 \stackrel{Eject}{\rightarrow} r_6\}$ which formalizes h_1,
- $F_{21} \stackrel{\text{def}}{=} \{r_8 \stackrel{Csends}{\rightarrow} r_{10}\}$ which formalizes h_2,
- $F_{22} \stackrel{\text{def}}{=} \{r_3 \stackrel{Rsends}{\rightarrow} r_5\}$ which formalizes h_3.

6.3 Example of the verification by parts

In this section, we present an application of partitioned model checking on the example of the protocol T=1. We use the tool SPIN [Hol91] for the verification by model checking. The test consists in splitting the transition system of the protocol T=1, and in choosing a PLTL property P verifiable by parts. Then, we verify P under fairness assumptions f on the global system, and on each part. Our goal is to illustrate that the property $f \Rightarrow P$ is verifiable in a partitioned way.

Figures 9 to 12 (see the appendix) represent the parts which are obtained by splitting the global transition system in Fig. 8 at the refined level from Definition 4.5.

Let us verify $P \stackrel{\text{def}}{=} \Box(CardF_2 = bl \Rightarrow \Diamond(CardF_2 = lb \wedge SenderF_2 = reader))$ which expresses that when a card sends a block, then it will inevitably send a last block. We verify P by refinement based parts under the fairness assumptions f_1', f_{21}, f_{22}. The PLTL formula which expresses the fairness assumptions is $f \stackrel{\text{def}}{=} f_1' \wedge f_{21} \wedge f_{22}$. The property to be verified is $Q \stackrel{\text{def}}{=} f \Rightarrow P$.

First, we checked that Q is satisfied on the global system of the T=1 protocol by the tool SPIN. Second, we verified Q on each part in the following way. Before checking Q on the part s_0 such as described in the Fig. 9, we simplified Q because the part s_0 is only concerned with the fairness assumptions f_{22}. This part does not contain a cycle, which is forbidden by the fairness assumptions f_1' or f_{21}, therefore these assumptions are useless to check Q in this part. Thus we simplified Q in $Q_1 \stackrel{\text{def}}{=} f_{22} \Rightarrow P$. We used SPIN to check that Q_1 is satisfied on the part s_0, therefore Q is satisfied on s_0. The part s_1 in Fig. 10 is concerned with the fairness assumptions f_{21}. So we simplified Q in $Q_2 \stackrel{\text{def}}{=} f_{21} \Rightarrow P$. The property Q_2 is satisfied on the part s_1, therefore the property Q is satisfied on s_1. The parts s_2 in Fig. 11 and s_3 in Fig. 12 do not contain cycles forbidden by the fairness assumptions, all their executions are fair. Thus they are not concerned with fairness assumptions. So, we simplified Q in $Q_3 \stackrel{\text{def}}{=} P$. The property Q_3 is satisfied on the parts s_2 and s_3,

REFINEMENT teg1ref **REFINES** teg1

SETS

$FRAME = \{bl, lb, ackb\}$

VARIABLES

$SenderF_2,\ Cstatus_2,\ CardF_2,\ ReaderF_2$

INVARIANT

$SenderF_2 \in SENDER \wedge CardF_2 \in FRAME \wedge ReaderF_2 \in FRAME \wedge Cstatus_2 = Cstatus_1 \wedge$
$\big(ReaderF_2 = bl \vee \big(\big(CardF_2 = ackb \vee CardF_2 = lb\big) \wedge SenderF_2 = reader\big)\big) \Leftrightarrow \big(Sender_1 = reader\big) \wedge$
$\big(CardF_2 = bl \vee \big(\big(ReaderF_2 = ackb \vee ReaderF_2 = lb\big) \wedge SenderF_2 = card\big)\big) \Leftrightarrow \big(Sender_1 = card\big)$

INITIALISATION

$SenderF_2 := reader \parallel Cstatus_2 := in \parallel CardF_2 := lb \parallel ReaderF_2 := lb$

EVENTS

Rsends $\widehat{=}$ **SELECT** $\big(SenderF_2 = reader \wedge Cstatus_2 = in$
$\wedge \big(CardF_2 = ackb \vee CardF_2 = lb\big)\big)$
THEN $SenderF_2 := card \parallel ReaderF_2 := lb$ **END**;

Csends $\widehat{=}$ **SELECT** $\big(SenderF_2 = card \wedge Cstatus_2 = in$
$\wedge \big(ReaderF_2 = ackb \vee ReaderF_2 = lb\big)\big)$
THEN $SenderF_2 := reader \parallel CardF_2 := lb$ **END**;

Eject $\widehat{=}$ **SELECT** $\big(\big(SenderF_2 = card \wedge ReaderF_2 = lb\big) \vee \big(SenderF_2 = reader \wedge CardF_2 = lb\big)\big) \wedge$
$Cstatus_2 = in$
THEN $Cstatus_2 := out$ **END**;

Cinsert $\widehat{=}$ **SELECT** $Cstatus_2 = out$
THEN $SenderF_2 := reader \parallel Cstatus_2 := in \parallel CardF_2 := lb \parallel ReaderF_2 := lb$ **END**;

Rblocksends $\widehat{=}$ **SELECT** $\big(SenderF_2 = reader \wedge Cstatus_2 = in \wedge \big(CardF_2 = ackb \vee CardF_2 = lb\big)\ \big)$
THEN $SenderF_2 := card \parallel ReaderF_2 := bl$ **END**;

Cblocksends $\widehat{=}$ **SELECT** $\big(SenderF_2 = card \wedge Cstatus_2 = in \wedge \big(ReaderF_2 = ackb \vee ReaderF_2 = lb\big)\ \big)$
THEN $SenderF_2 := reader \parallel CardF_2 := bl$ **END**;

Racksends $\widehat{=}$ **SELECT** $\big(SenderF_2 = reader \wedge Cstatus_2 = in \wedge CardF_2 = bl\big)$
THEN $SenderF_2 := card \parallel ReaderF_2 := ackb$ **END**;

Cacksends $\widehat{=}$ **SELECT** $\big(SenderF_2 = card \wedge Cstatus_2 = in \wedge ReaderF_2 = bl\big)$
THEN $SenderF_2 := reader \parallel CardF_2 := ackb$ **END**;

FAIRNESS$=\{Eject\ ,\ Csends\ if\ (CardF_2 = bl),\ Rsends\ if\ (ReaderF_2 = bl)\}$;

END

Figure 7: Refined B specification of the protocol T=1 under fairness assumptions

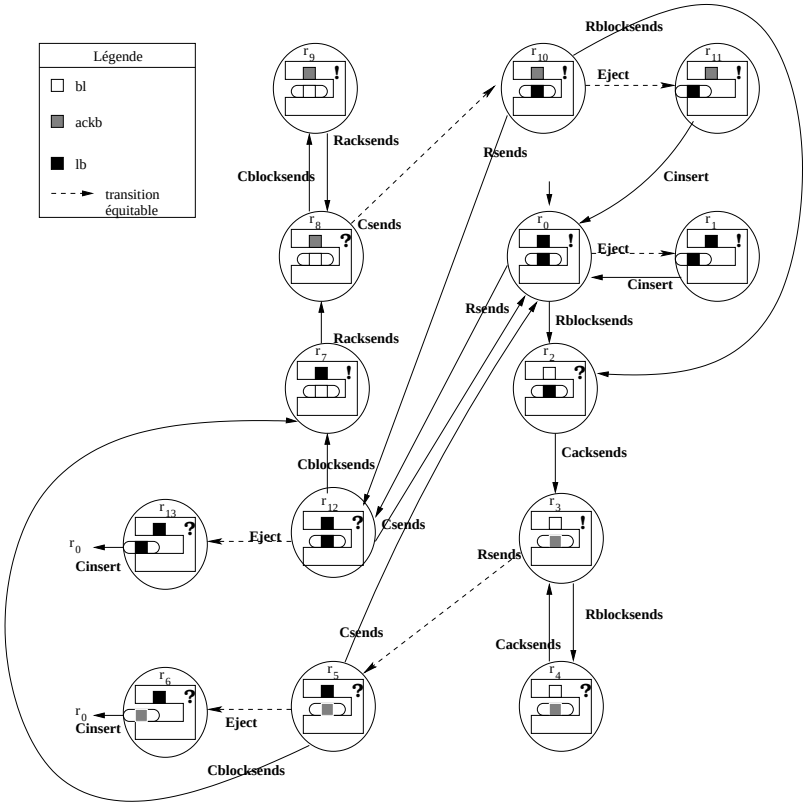

Figure 8: Refined fair transition system of the protocol T=1

therefore the property Q is satisfied also on these parts. Accordingly to Theorem 5.2, the property Q is verifiable by refinement based parts (see Definition 4.5). Since all the parts satisfy Q, the global system of the protocol T=1 satisfies Q.

6.4 Counter example if the condition of decomposition by refinement does not holds

In this section, we give an example that shows that the approach of verifying by parts PLTL property under fairness assumptions is not correct, when the condition (see Definition 4.5) of the decomposition by refinement is not satisfied. We propose a decomposition of the system so that the executions of the parts cannot leave the τ-cycles and fair exiting cycles, contained in the parts resulting from the decomposition. In this example, we note $(r_i \stackrel{a_i}{\rightarrow} \ldots \stackrel{a_{i+n}}{\rightarrow} r_i)^*$ a finite fragment of an execution which runs around a cycle finitely many times. We note $(r_i \stackrel{a_i}{\rightarrow} \ldots \stackrel{a_{i+n}}{\rightarrow} r_i)^\omega$ an infinite fragment of an execution which runs around a cycle infinitely many times.

Figures 13 and 14 represent the parts s'_1 and s'_2 obtained by splitting the global system of the protocol T=1, without using the refinement. Notice that, we have not used a refinement based partitioning but the partition of the transitions of the system.

Let us verify the PLTL property $P' \stackrel{\text{def}}{=} \Box(SenderF_2 = reader \ \wedge \ ReaderF_2 = lb \ \wedge \ Cstatus_2 = in \Rightarrow \Diamond(ReaderF_2 = lb \wedge Cstatus_2 = out))$, which expresses that if the reader sends a last block then the card eventually will eject. We verify P' under the fairness assumptions $f = f'_1 \ \wedge \ f_{21} \ \wedge \ f_{22}$. Notice that P' is verifiable by parts, because the Büchi automaton of $\neg P'$ belongs to the class $\mathcal{B}_{mod}$. The property to be verified by parts is $Q' \stackrel{\text{def}}{=} f \Rightarrow P'$.

We verified Q' on the global system of the protocol T=1 using SPIN. Then, we verified Q' on each part, s'_1 and s'_2. The results are that Q' is violated on the global system and satisfied on the parts. In the part s'_2, the executions are not fair. Since f is false on s'_2 then $Q' \stackrel{\text{def}}{=} f \Rightarrow P'$ is satisfied on s'_2. Q' is satisfied on the part s'_1 because the only computations of this part are those which reach infinitely many times the states r_0 and r_1 in Fig. 13. These executions satisfy Q'. The other executions of s'_1 are not fair, therefore they satisfy Q'. Consequently, the property Q' is not verifiable by parts although the Büchi automaton of $\neg P'$ belongs to the class $\mathcal{B}_{mod}$, i.e. is verifiable by parts. We recall that a PLTL property P under fairness assumptions $f \Rightarrow P$ is verifiable by parts under the following condition : "if the property is not satisfied on the global system, then there is a part which violates the property".

With such a decomposition of the global system, the method of verifying by parts is not correct because the parts do not contain fragments of some computations of the global system. For example, the computation

$$\sigma_1 \stackrel{\text{def}}{=} (r_0 \stackrel{Rblocksends}{\rightarrow} r_2 \stackrel{Cacksends}{\rightarrow} (r_3 \stackrel{Rblocksends}{\rightarrow} r_4 \stackrel{Cacksends}{\rightarrow} r_3)^* \stackrel{Rsends}{\rightarrow} r_5 \stackrel{Eject}{\rightarrow} r_6 \stackrel{Cinsert}{\rightarrow} r_0)^\omega$$

in Fig. 8, does not have fragments in the parts s'_1 and s'_2. Indeed the executions

$$\sigma_2 \stackrel{\text{def}}{=} r_0 \stackrel{Rblocksends}{\rightarrow} r_2 \stackrel{Cacksends}{\rightarrow} (r_3 \stackrel{Rblocksends}{\rightarrow} r_4 \stackrel{Cacksends}{\rightarrow} r_3)^\omega$$

in the part s'_1 and

$$\sigma_3 \stackrel{\text{def}}{=} (r_0 \stackrel{Rsends}{\rightarrow} r_{12} \stackrel{Csends}{\rightarrow} r_0)^\omega$$

in the part s'_2, are not extensions of fragments of the execution σ_1. As σ_2 and σ_3 are not fair, then they satisfy the property Q'. The method fails because there are computations of the global transition system which are not verified in the parts. This because their fragments do not exist in the computations of the parts. On the other hand the decomposition by refinement ensures that all the fragments of all the computations of the global system are in the parts. For example, the computation σ_1 is broken into two fragments which we find in the computations of the parts obtained by refinement. The first fragment of σ_1 is

$$\varphi_1 \stackrel{\text{def}}{=} r_0 \stackrel{Rblocksends}{\rightarrow} r_2 \stackrel{Cacksends}{\rightarrow} (r_3 \stackrel{Rblocksends}{\rightarrow} r_4 \stackrel{Cacksends}{\rightarrow} r_3)* \stackrel{Rsends}{\rightarrow} r_5 \stackrel{Eject}{\rightarrow} r_6$$

which is a prefix of the computation

$$\sigma_4 \stackrel{\text{def}}{=} r_0 \stackrel{Rblocksends}{\rightarrow} r_2 \stackrel{Cacksends}{\rightarrow} r_3 \stackrel{Rblocksends}{\rightarrow} (r_4 \stackrel{Cacksends}{\rightarrow} r_3 \stackrel{Rblocksends}{\rightarrow} r_4)^* \stackrel{Rsends}{\rightarrow} r_5 \stackrel{Eject}{\rightarrow} (r_6 \stackrel{skip}{\rightarrow} r_6)^\omega$$

in the part s_0 in Fig. 9. We find also the second fragment

$$\varphi_2 \stackrel{\text{def}}{=} r_6 \stackrel{Cinsert}{\rightarrow} r_0$$

as a prefix of the computation

$$\sigma_5 \stackrel{\text{def}}{=} r_6 \stackrel{Cinsert}{\rightarrow} r_0 \stackrel{Eject}{\rightarrow} (r_1 \stackrel{skip}{\rightarrow} r_1)^\omega$$

in the part s_2 in Fig. 11. Let us note that σ_1 is the concatenation of these two fragments. The two executions σ_4 and σ_5 which contain the fragments φ_1 and φ_2 are fair, therefore they satisfy f. Since σ_4 and σ_5 satisfy also $Q \stackrel{\text{def}}{=} f \Rightarrow P$ as verified previously – Q is the property verified by parts with a decomposition by refinement in Section 6.3 – therefore they satisfy P. Since P is verifiable by parts, then the proof that P is satisfied on σ_4 and σ_5 is a proof that P is satisfied on σ_1.

7 Performance of the approach of the verification by part

In this section, we show the experimental results of the approach of the verification by refinement based parts. The expected performance is the capacity of this approach to verify a large set of different type of properties at the refined level. So, we verify two examples of applications. The first example is the protocol T=1, the second one is the car wind-screen wipers system. For each example we choose different PLTL properties to verify. The properties must express the general behaviors of the systems.

At the end of this section, we compare our approach with the Pnueli's approach [KPR98, KPRS01], that is another verification approach under fairness assumptions.

The protocol T=1 at the refined level is described in the last section. We propose to verify by parts the following PLTL properties which express the main behaviors of the protocol :

- These properties express that always when a device send a block then it will inevitably send a last block:

  $P_1 \stackrel{\text{def}}{=} \Box(CardF_2 = bl \Rightarrow \Diamond(CardF_2 = lb))$. $P_2 \stackrel{\text{def}}{=} \Box(ReaderF_2 = bl \Rightarrow \Diamond(ReaderF_2 = lb))$.

- These properties express that always when a device send a block then the other device will inevitably send an acknowledgment block:

  $P_3 \stackrel{\text{def}}{=} \Box(CardF_2 = bl \Rightarrow \Diamond(ReaderF_2 = ackb))$. $P_4 \stackrel{\text{def}}{=} \Box(ReaderF_2 = bl \Rightarrow \Diamond(CardF_2 = ackb))$.

- This property expresses that when the card sends a block and the reader sends an acknowledgment block, the card and the reader will respectively send an acknowledgment block and a block:

  $P_5 \stackrel{\text{def}}{=} \Box(CardF_2 = bl \ \wedge \ ReaderF_2 = ackb \Rightarrow \Diamond(CardF_2 = ackb \ \wedge \ ReaderF_2 = bl))$.

- This property expresses the alternation of sending of the messages between the card and the reader:

  $P_6 \stackrel{\text{def}}{=} \Box(SenderF_2 = card \Rightarrow \Diamond(SenderF_2 = reader))$.

The second application verified is the car wind-screen wipers system. At the refined level the system is composed of a control level, a rain sensor and two (left and right) wind-screen wipers. The control level can select the mode - automatic or manual- of the wiper system. The left and the right wipers have the same behavior. The rain sensor can detect the rain amount (no rain, small rain, strong rain). We have verified six properties on this application. The results are shown in the following section.

7.1 Results of the verification

The following table indicates the results. We give the number of properties to verify, how many are globally true, how many globally false, and how many have been successfully verified by refinement based parts.

Example	Properties	Globally true	Globally false	Verified by parts
protocol T=1	6	5	1	4
wind-screen wipers	6	6	0	4

These results show that we have successfully verified four properties by parts amongst the six that were expressed, on the protocol T=1 system and the car wind-screen wipers system. In the case of the protocol T=1, the verification failed for properties P_5 and P_6. As P_5 is globally false, then there is at least a part on which it is false, that is s_0. In contrast the property P_6 is globally true and our method failed to prove it by refinement based parts (it is false on parts s_0 and s_1). This is due to the fact that P_6 is not a *new* property. It expresses an abstract behavior of the system, and should have been verified by parts at a former level of refinement. Properties P_1, P_2, P_3 and P_4 express new behaviors of the system at the refined level, and as expected, they have been successfully verified by parts.

We obtained the same results for the verification of the car wind-screen wipers system. The approach of the verification by parts fails in verifying two properties that have been verified at the abstract level. However, the approach succeeds in verifying properties that express the new behaviors of the system at the refined level.

These results show that as it would be interesting to study the relation between our approach of verification, the new properties at the refined level, and the abstract properties in order to characterize the properties for which a verification by parts is suited.

7.2 Comparison with Pnueli's approach of verification

An interesting method was proposed in [KPR98, KPRS01], this is a symbolic model checking of PLTL properties under fairness assumptions. This approach removes the fairness assumptions from the formula to verify. It deals with a fairness assumptions at the algorithmic level instead of specifying them as of a part of the formula to be verified. Fairness assumptions are expressed as Büchi (for weak fairness) and Street (for strong fairness) automata acceptance conditions. So, this algorithm verifies the property P instead of verifying $f \Rightarrow P$. The verification consists of the emptiness checking which is implemented with BDD.

This approach treats the problem of the combinatorial explosion of model checking by simplifying the formula to verify under fairness approach. In our case we treats the problem of the combinatorial explosion by the following way. We partition the transition system and we verify a property separately on each part. Another difference is that we adapted the automata algorithm of Vardi and Wolper [VW86] instead of a symbolic algorithm. However, we can combine our approach with Pnueli's approach. So in order to verify a property P under fairness assumptions f, we verify the simplified formula P by the partitioning way. So we will use Pnueli's approach to express fairness assumptions by Street and Büchi automata conditions on each transition system of the parts. Then we will exploit our approach to split the global transition system using the fair refinement relation. Finally we will verify P on each part using symbolic model checking instead of verifying $f \Rightarrow P$ on the global system.

8 Conclusion and Future Works

In this paper, we extend the partitioned model checking technique presented in [JMM01] to handle the fairness constraints of the system environment. Our goal is to verify the PLTL properties under fairness assumptions by part. When the fairness constraints of the environment are expressed by fairness assumptions, the verification by model checking of a partitioned property P under fairness assumptions f supposes verifying

by model checking the new property $Q \stackrel{\text{def}}{=} f \Rightarrow P$ on the transition system. However, the property Q does not necessarily belong to the class of properties verifiable by part.

Our contribution in this paper is to prove that the split of the transition system into parts, using the fair refinement relation, makes the property Q verifiable by refinement based parts when P is verifiable by part. The use of the fair refinement to split a transition system allows us to obtain refinement based parts which contains computations. This is a sufficient condition to verify by refinement based parts the property $Q \stackrel{\text{def}}{=} f \Rightarrow P$. To handle the fairness constraints, we have proposed to use a fair transition system to model a reactive system and its fairness environment. This framework is a transition system which contains only computations.

The complexity of the refinement verification is linear in the size of the refined system, because it necessitates only an enumeration of the refined model when the gluing relation is a function. Therefore the verification by parts is interesting because the additive decomposition of a system comes for free from the refinement verification.

In the future, we plan to implement the partitioned model checking technique so that we can evaluate its performance on industrial applications. As we saw in the example of the protocol T=1, we can simplify the fairness assumptions on each part because all the assumptions do not concern every part. So, the simplification process needs to be further studied. We plan also to combine the approach of model checking by part to the approach [CJ03] of model checking under fairness assumptions which exploits the fair refinement in order to reduce the size of the formula to be verified under fairness assumptions.

Also, we must give simplification rules to translate fairness assumptions, as expressed in the event systems, into PLTL formulae, as used by the usual model checking algorithms. We plan to study a variation of this method which does not requires the condition (b) in Definition 2.8. In this case the fairness assumptions are not expressible in PLTL, but in a logic of actions as the μ-calculus [CGP99].

References

[Abr96] J.-R. Abrial. Extending B without changing it (for developing distributed systems). In *1st Conference on the B method*, pages 169–190, Nantes, France, November 1996.

[AdAHM99] R. Alur, L. de Alfaro, T. A. Henzinger, and F. Y. C. Mang. Automating modular verification. In *Proc. of the 10th Conf. on Concurrency Theory*, volume 1664 of *Lecture Notes in Computer Science*, pages 82–97. Springer-Verlag, July 1999.

[AM98] J.-R. Abrial and L. Mussat. Introducing dynamic constraints in B. In *2nd Conference on the B method, France*, volume 1393 of *Lecture Notes in Computer Science*, pages 83–128. Springer Verlag, April 1998.

[BCJK01] F. Bellegarde, S. Chouali, J. Julliand, and O. Kouchnarenko. Comment limiter la spécification de l'équité dans les systèmes d'événements B. In *Approches Formelles dans l'Assistance au Développement de Logiciels (AFADL'01)*, pages 205–220, Nancy, France, 2001.

[BCMD90] J. R. Burch, E. M. Clarke, K. L. McMillan, and L. J. Dill, D. L.and Hwang. Symbolic model checking : 10^{20} states and beyond. In *Proc. of the 5th Annual IEEE Symposium on Logic in Computer Science*, pages 428–439, Philadelphia, Pensylvania, USA, June 1990. IEEE Computer Society Press.

[BJK00] F. Bellegarde, J. Julliand, and O. Kouchnarenko. Ready-simulation is not ready to express a modular refinement relation. In *Proc. Int. Conf. on Fondamental Aspects of Software Engineering, FASE'2000*, volume 1783 of *Lecture Notes in Computer Science*, pages 266–283. Springer-Verlag, April 2000.

[BJK02] F. Bellegarde, J. Julliand, and O. Kouchnarenko. Synchronized parallel composition of event systems in B. In Didier Bert, Jonathan P. Bowen, Martin C. Henson, and Ken Robinson,

editors, *Formal Specification and Development in Z and B, proc. of the 2nd Int. Conf. of B and Z Users, (ZB 2002)*, volume 2272 of *LNCS*, pages 436–457, Grenoble, France, January 2002. Springer.

[CC77] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *4th ACM Symp. on Principles of Programming Languages (POPL'77)*, pages 238–252, Los Angeles, California, USA, 1977. ACM Press.

[CES86] E. M. Clarke, E. A. Emerson, and P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specification. In *ACM transactions on Programming Languages and Systems*, volume 2, pages 244–263, 1986.

[CGK97] S.C. Cheung, D. Giannakopoulou, and J. Kramer. Verification of liveness properties using compositional reachability analysis. In *ESEC/SIGSOFT FSE, 6th European Software Engineering Conference*, volume 1301 of *Lecture Notes in Computer Science*, pages 227–243. Springer-Verlag, November 1997.

[CGL94] E. M. Clarke, O. Grumberg, and D. E. Long. Model checking and abstraction. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 16(5):1512–1542, 1994.

[CGP99] E. M. Jr Clarke, O Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 1999.

[Cho03] S. Chouali. *Contribution du raffinement à la vérification de systèmes sous hypothèses d'équité*. PhD thesis, Université de Franche-Comté, Besançon, France, December 2003.

[CJ03] S. Chouali and J. Julliand. Model checking des propriétés dynamiques sous hypothèses d'équité, exploitant le raffinement. In *Approches Formelles dans l'Assistance au Développement de Logiciels (AFADL'03)*, pages 277–291, Rennes, France, 2003.

[CLM89] E. M. Clarke, D. E. Long, and K. L. McMillan. Compositional Model Checking. In *Proc. of the 4th Annual Symposium on Logic in Computer Science*, pages 353–362, Pacific Grove, California, USA, 1989. IEEE Computer Society Press.

[DF95] J. Dingel and T. Filkorn. Model checking for infinite state systems using data abstraction, assumption-commitment style reasoning and theorem proving. In *Computer Aided Verification, CAV'95*, volume 939 of *Lecture Notes in Computer Science*, pages 54–69, Liège, Belgique, July 1995. Springer-Verlag.

[DHWT91] D.L. Dill, A.J. Hu, and H. Wong-Toi. Checking for language inclusion using simulation relation. In *proc. of Computer Aided Verification (CAV'91)*, volume 575 of *LNCS*, pages 255–265, 1991.

[DJK03] C. Darlot, J. Julliand, and O. Kouchnarenko. Refinement preserves PLTL properties. In *3rd Int. Conf. of B and Z Users, ZB'03*, volume 2651 of *LNCS*, pages 408–420, Turku, Finland, September 2003.

[edNE92] Comité européen de Normalisation. En27816-3. European standard - identification cards - integrated circuit(s) card with contacts - electronic signal and transmission protocols. Technical Report ISO/CEI 7816-3, ISO, 1992.

[GL94] O. Grumberg and David E. Long. Model checking and modular verification. *ACM Transactions on Programming Languages and Systems*, 16(3):843–871, May 1994.

[HKR97] T.A. Henzinger, O. Kupferman, and S. Rajamani. Fair simulation. In *Proc. of the 8th Conference on Concurrency Theory*, volume 1243 of *LNCS*, pages 273–287, Warsaw, Poland, July 1997. Springer-Verlag.

[Hol91] G. Holzmann. *Design and validation of computer protocols*. Prentice hall software series, 1991.

[JMM01] J. Julliand, P.-A. Masson, and H. Mountassir. Vérification par model-checking modulaire des propriétés dynamiques introduites en B. *Technique et Science Informatiques*, 20(7):927–957, September 2001.

[KKPR99] Y. Kesten, A. Klein, A. Pnueli, and G. Raanan. A perfecto verification: combining model checking with deductive analysis to verify real-life software. In *Formal methods'99, Toulouse, France*, volume 1708 of *Lecture Notes in Computer Science*, pages 173–194. Springer-Verlag, October 1999.

[KL93] R. P. Kurshan and L. Lamport. Verification of a multiplier: 64-bits and beyond. In *Proc. of the 5th Int. Conf. on Computer Aided Verification (CAV'93)*, volume 697 of *LNCS*, pages 166–179, Elounda, Greece, 1993. Springer-Verlag.

[KLM+02] R. Kurshan, V. Levin, M. Minea, D. Peled, and H. Yenign. Combining software and hardware verification techniques. *Formal Methods in System Design*, 21:251–280, 2002.

[KP88] S. Katz and D. A. Peled. An efficient verification method for parallel and distributed programs. In *Workshop on Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency*, volume 354 of *LNCS*, pages 489–507. Springer-Verlag, 1988.

[KPR98] Y. Kesten, A. Pnueli, and L. Raviv. Algorithmic verification of linear temporal logic specifications. In *Proc. of the 25th Int. Colloquium on Automata, Languages, and Programming (ICALP 1998)*, volume 1443 of *LNCS*, pages 1–16. Springer-Verlag, 1998.

[KPRS01] Y. Kesten, A. Pnueli, L. Raviv, and E. Shahar. Model checking with strong fairness. Technical Report MCS01-07, Weizmann Institute of Science, March 2001.

[KV96] O. Kupferman and M.Y. Vardi. Verification of fair transition systems. In *Proc. 8th Int. Conf. on Computer Aided Verification (CAV'96)*, volume 1102 of *LNCS*, pages 372–382. Springer-Verlag, 1996.

[KV98] O. Kupferman and M. Y. Vardi. Modular model checking. In *Compositionality Workshop, COMPOS'97*, volume 1536 of *Lecture Notes in Computer Science*, pages 381–401. Springer-Verlag, 1998.

[LP85] O. Lichtenstein and A. Pnueli. Checking that finite state concurrent programs satisfy their linear specification. In *Proc. of the 12-th ACM Symp. on Principles of Programming Languages (POPL'85)*, pages 97–107, New Orleans, USA, 1985.

[Mas01] P.-A. Masson. *Vérification par model-checking modulaire de propriétés dynamiques PLTL exprimées dans le cadre de spécifications B événementielles*. PhD thesis, Université de Franche-Comté, December 2001.

[McM93] Ken McMillan. *Symbolic model-checking*. Kluwer Academic Publishers, 1993.

[MMJ00] P.-A. Masson, H. Mountassir, and J. Julliand. Modular verification for a class of PLTL properties. In *Proc. of the 2nd Int. Conf. on Integrated Formal Methods (IFM 2000)*, volume 1945 of *LNCS*, pages 398–419, Dagstuhl Castle, Germany, november 2000. Springer-Verlag.

[MP92] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag - ISBN 0-387-97664-7, 1992.

[MP95] Z. Manna and A. Pnueli. *Temporal Verification of Reactive Systems :Safety*. Springer-Verlag - ISBN 0-387-94459-1, 1995.

[QS82] J.-P. Queille and J. Sifakis. Specification and verification of concurrent systems in CESAR. In *Proc. of the Int. Symp. on Programming*, volume 137 of *LNCS*, pages 337–351, Turin, Italy, April 1982. Springer-Verlag.

[Tol99] Silvan Toledo. A survey of out-of-core algorithms in numerical linear algebra. *External Memory Algorithms and Visualization*, 1999.

[VW86] M. Y. Vardi and P. Wolper. An automata theoretic approach to automatic program verification. In *Proc. of the 1st IEEE Symp. on Logic in Computer Science (LICS'86)*, pages 332–344, Cambridge, USA, June 1986.

[WG93] P. Wolper and P. Godefroid. Partial-order methods for temporal verification. In *Proc. of the Int. Conf. on Concurrency Theory (CONCUR'93)*, volume 715 of *LNCS*, pages 233–246, Hildesheim, August 1993. Springer-Verlag.

Appendix

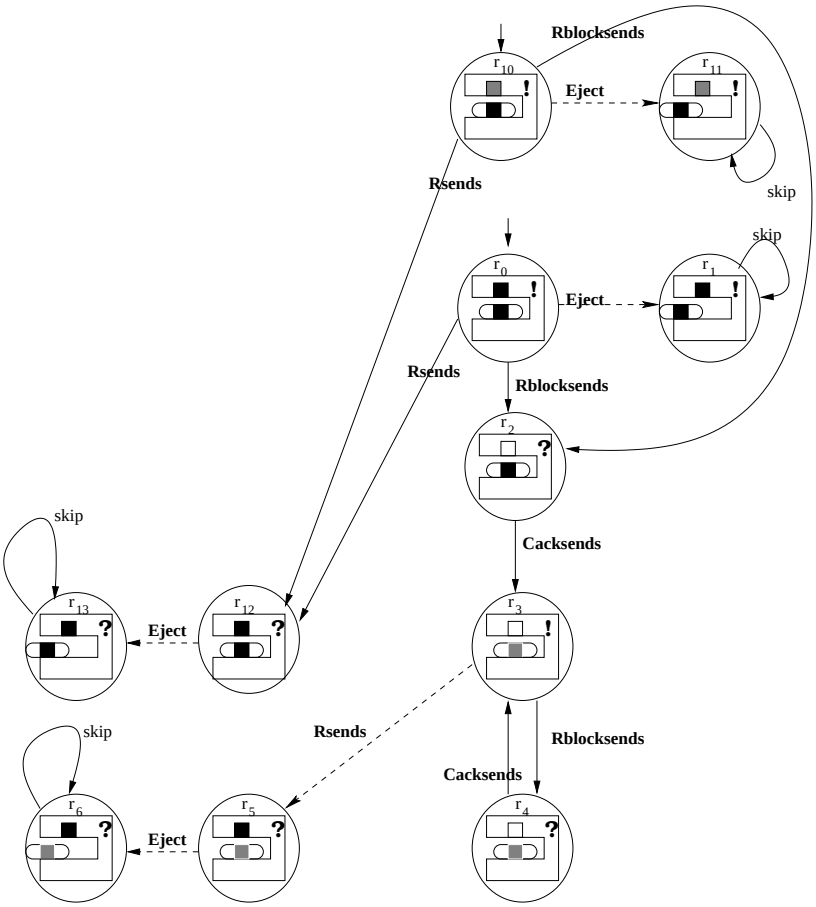

Figure 9: The part s_0 of the protocol T=1

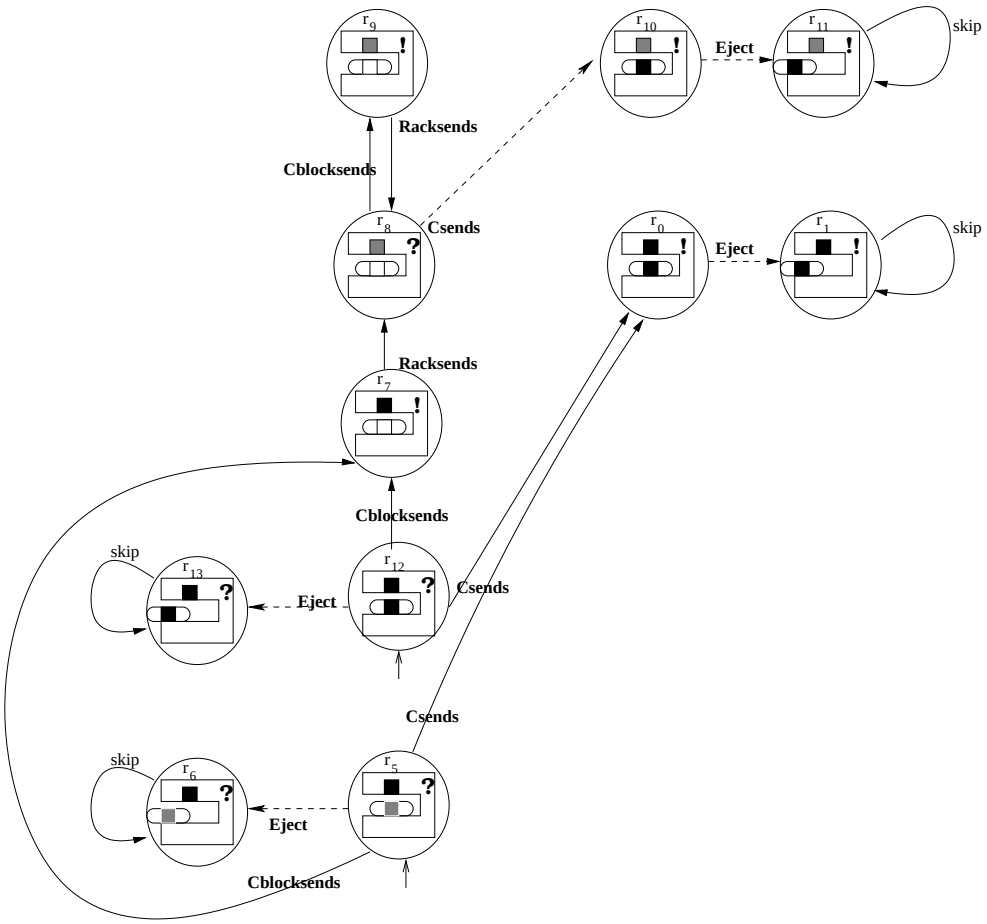

Figure 10: The part s_1 of the protocol T=1

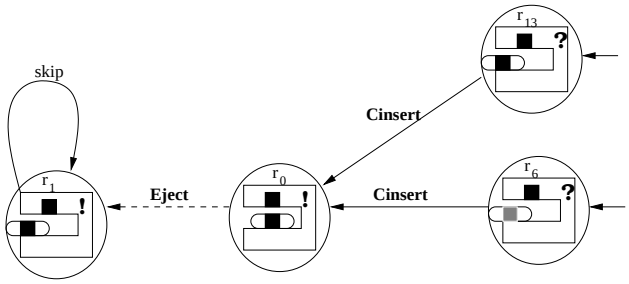

Figure 11: The part s_2 of the protocol T=1

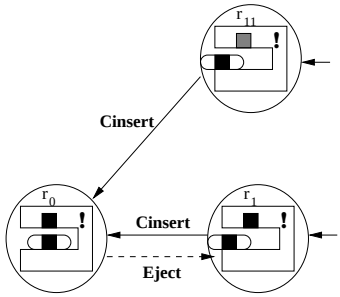

Figure 12: The part s_3 of the protocol T=1

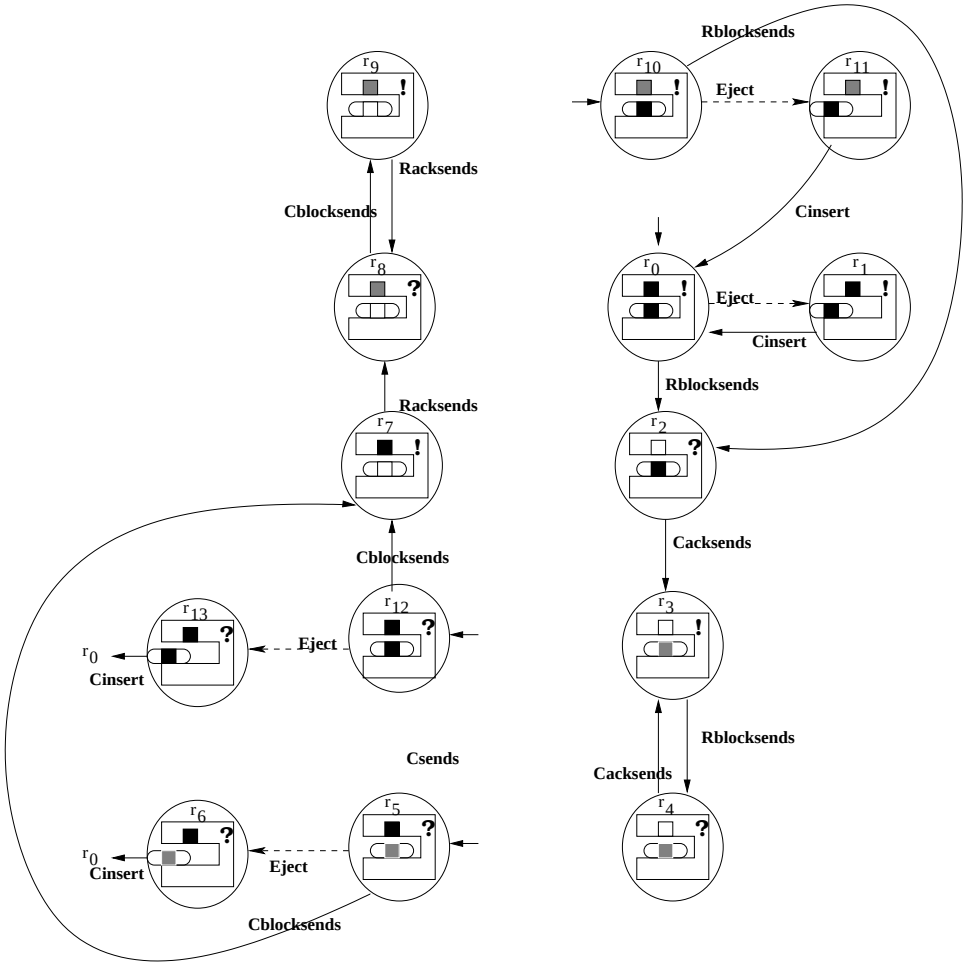

Figure 13: The part s_1' of the protocol T=1

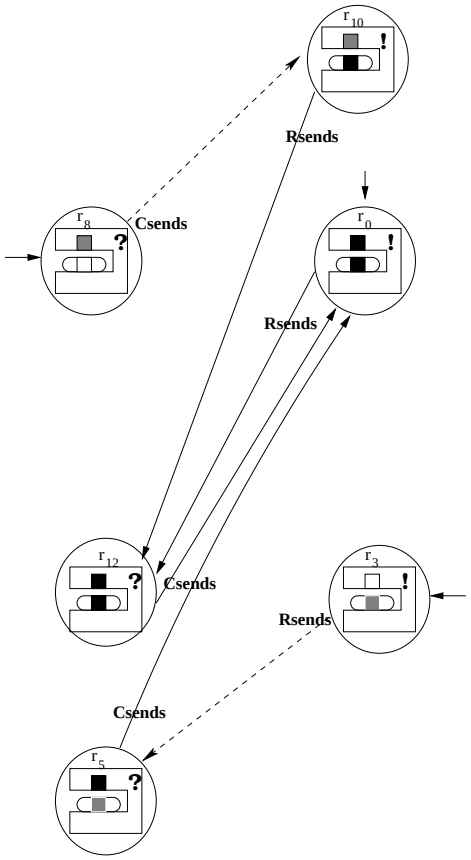

Figure 14: The part s_2' of the protocol T=1